

# Gas Clumping in Self-Consistent Reionisation Models

Kristian Finlator<sup>1†</sup>, S. Peng Oh<sup>1</sup>, Feryal Özel<sup>2</sup>, Romeel Davé<sup>2</sup>

<sup>1</sup> Department of Physics, University of California, Santa Barbara, CA 93106, USA

<sup>†</sup> Hubble Fellow

<sup>2</sup> Astronomy Department, University of Arizona, Tucson, AZ 85721, USA

13 September 2012

## ABSTRACT

We use a suite of cosmological hydrodynamic simulations including a self-consistent treatment for inhomogeneous reionisation to study the impact of galactic outflows and photoionisation heating on the volume-averaged recombination rate of the intergalactic medium (IGM). By incorporating an evolving ionising escape fraction and a treatment for self-shielding within Lyman limit systems, we have run the first simulations of “photon-starved” reionisation scenarios that simultaneously reproduce observations of the abundance of galaxies, the optical depth to electron scattering of cosmic microwave background photons  $\tau_{\text{es}}$ , and the effective optical depth to Ly $\alpha$  absorption at  $z = 5$ . We confirm that an ionising background reduces the clumping factor  $C$  by more than 50% by smoothing moderately-overdense ( $\Delta = 1\text{--}100$ ) regions. Meanwhile, outflows increase clumping only modestly. The clumping factor of ionised gas is much lower than the overall baryonic clumping factor because the most overdense gas is self-shielded. Photoionisation heating further suppresses recombinations if reionisation heats gas above the canonical 10,000 K. Accounting for both effects within our most realistic simulation,  $C$  rises from  $< 1$  at  $z > 10$  to 3.3 at  $z = 6$ . We show that incorporating temperature- and ionisation-corrected clumping factors into an analytical reionisation model reproduces the numerical simulation’s  $\tau_{\text{es}}$  to within 10%. Finally, we explore how many ionising photons are absorbed during the process of heating filaments by considering the overall photon cost of reionisation in analytical models that assume that the IGM is heated at different redshifts. For reionisation redshifts of 9–10, cold filaments boost the reionisation photon budget by  $\sim 1$  photon per hydrogen atom.

**Key words:** cosmology: theory — radiative transfer — hydrodynamics — methods: numerical

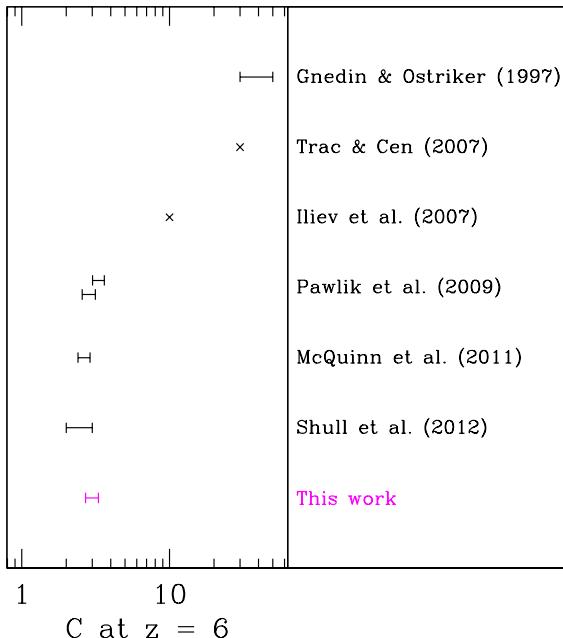
## 1 INTRODUCTION

Much of the current interest in cosmological hydrogen reionisation may be distilled to two related questions: (1) When did it occur?; and (2) What sources dominated the ionising photon budget? Theoretical efforts to address the first question have traditionally considered two observational constraints: the optical depth to Thomson scattering of cosmic microwave background (CMB) photons,  $\tau_{\text{es}}$ , and the opacity of the Lyman- $\alpha$  forest (Trac & Gnedin 2009, and references therein). While these works have shed much light on the nature of inhomogeneous reionisation, their predictions have until recently been uncertain owing to the unknown relation between ionising luminosity and halo mass (McQuinn et al. 2007). With the recent installation of the Wide-Field Camera 3 on board the Hubble Space Telescope, a new generation of observations is now constraining the abundance and colors of galaxies back to  $z = 8$  and beyond (Finkelstein et al. 2010; Bouwens et al. 2011;

Dunlop et al. 2012; González et al. 2011; Grazian et al. 2011; McLure et al. 2011; Oesch et al. 2012). These observations may be used to constrain the relationship between star formation rate and halo mass via abundance-matching studies (Trenti et al. 2010; Muñoz & Loeb 2011). Furthermore, the relationship between ionising luminosity and star formation rate is now constrained by the direct detection of escaping Lyman continuum flux from high-redshift galaxies (Shapley et al. 2006; Siana et al. 2010; Nestor et al. 2011). Hence there are now tenuous observational links connecting halo mass with ionising luminosity, a major step toward understanding the relationship between high-redshift galaxies and their environment.

By combining observations of galaxies, the CMB, and the Lyman- $\alpha$  forest, it is possible to ask whether the observed ionising sources and reionisation history are compatible (for example, Haardt & Madau 2012; Kuhlen & Faucher-Giguère 2012). Two uncertainties hinder

## 2 Gas Clumping at $z \geq 5$



**Figure 1.** Our ionisation- and temperature-corrected IGM clumping factor at  $z = 6$  in comparison to previous calculations. Our work indicates that the clumping factor lies within the range 2.7–3.3, which is consistent with the low values that have been found in recent works.

progress toward answering this question. The first remains the fraction of ionising photons  $f_{\text{esc}}$  that escape into the IGM: direct constraints are only available at lower redshifts and higher luminosities than are relevant to reionisation. Furthermore,  $f_{\text{esc}}$  is sensitive to parsec-scale physical processes that occur deep within galaxies' interstellar media and is difficult to model numerically (Fernandez & Shull 2011, and references therein). The second uncertainty involves the recombination rate of the reionisation-epoch IGM. The unknown ratio of the IGM's true recombination rate to its hypothetical rate under the assumption of uniform density and temperature is often referred to as the IGM clumping factor  $C$ . We will define  $C$  explicitly in Section 3; for the present the important point is that it is proportional to the volume-averaged hydrogen recombination rate in the IGM. High values of  $C$  require abundant faint galaxies and/or a large  $f_{\text{esc}}$  while low values of  $C$  are easier to reconcile with existing observations.  $C$  is sensitive to the IGM's density, ionisation, and temperature fields, hence an accurate estimate requires three-dimensional numerical simulations.

To this end,  $C$  has been studied numerous times using numerical simulations. We compile predictions of the clumping factor at  $z = 6$  from a variety of works in Figure 1. The early work of Gnedin & Ostriker (1997) yielded a clumping factor of 30–50 (their Figure 2) depending on the choice of definition. However, this clumping factor included the recombinations that occur in gas whose ionisation state is dominated by the local field rather than by the global extragalactic ultraviolet ionising background (EUVB), hence it is an overestimate of the recombination rate within the gas that is conventionally associated with the IGM (see also Gnedin 2000 and McQuinn et al. 2011). For the same

reason, Trac & Cen (2007) also inferred a clumping factor of roughly 30 from simulations that subtended a significantly larger cosmological volume.

Iliev et al. (2007) applied this insight to calculate an improved clumping factor using high-resolution N-body simulations whose spatial resolution was comparable to the Jeans length in the IGM. By excluding matter within virialized halos, they derived a redshift-dependent fitting function that climbs to  $\approx 10$  at  $z = 6$ , a value which was also found in the recent N-body simulations of Raičević & Theuns (2011). This work confirmed that physically-motivated density cuts reduce  $C$ . However, as the authors noted, their calculation did not account for the fact that photoionisation heating (photoheating) tends to smooth the diffuse IGM. We will confirm in Figure 5 that this leads to an overestimated clumping factor.

Pawlik et al. (2009) corrected this problem by running cosmological hydrodynamic simulations that model reionisation using a spatially-homogeneous EUVB. They used several plausible density cuts to calculate the IGM clumping factor as a function of redshift. We show their predictions assuming that the IGM corresponds to overdensity cuts of  $\Delta \equiv \rho/\langle \rho \rangle < 50$  (lower range) and  $\Delta < 100$  (upper range). This work generally confirmed the impression that considering only low-density gas and accounting for the tendency for the photoheating of diffuse regions both reduce  $C$ . However, it suffered from two drawbacks. First, their use of a spatially-homogeneous EUVB pressurized regions that should in reality be self-shielded, leading to uncertainty in the gas density distribution from which the predicted clumping factors were derived. Second, their analysis did not account explicitly for the IGM temperature, which could suppress the recombination rate further.

In a qualitatively similar work, Shull et al. (2012) ran a cosmological hydrodynamic simulation with a spatially-homogeneous EUVB and computed the temperature-corrected clumping factor within gas that satisfied cuts on baryon overdensity, temperature, metallicity, and ionisation state. Although these extra cuts render direct comparisons with Pawlik et al. (2009) uncertain, the inferred range is consistent. Both works are also in reasonable agreement with Faucher-Giguère et al. (2008), who used the gas density distribution of Miralda-Escudé et al. (2000) to find that  $C = 3\text{--}4$  if the clumping factor is averaged only over gas with overdensity  $\Delta$  less than 67.

The uncertainty in the derived clumping factor may be reduced by using radiation transport simulations to identify directly the gas that is self-shielded. McQuinn et al. (2011) use this approach in post-processing to derive a clumping factor of 2.4–2.9. Their calculations do not account for the IGM temperature and, like the works of Pawlik et al. (2009) and Shull et al. (2012), begin with hydrodynamic simulations in which the moderately-overdense gas is incorrectly exposed to the EUVB (although they confirm directly that shielding gas that is dense enough to form stars has no impact). Nonetheless, their results are clearly consistent with a relatively low clumping factor.

The goal of the present work is to combine all of these ideas into a single radiation hydrodynamic framework that accounts realistically for the time and spatial dependence of the IGM's density, ionisation, and temperature fields. Our fiducial simulation reproduces the observed UV luminosity

function of galaxies at  $z \geq 6$  (Finlator et al. 2011) and is marginally consistent with the effective optical depth to absorption of Ly $\alpha$  photons and  $\tau_{\text{es}}$ . Hence, our estimated IGM recombination rate is constrained by a variety of observations. As Figure 1 indicates, we will confirm that effects that have been neglected in recent works based on hydrodynamic simulations such as incorrectly heating dense regions and ignoring the gas temperature are indeed small. Of course, the purpose of clumping factors is to inform analytical models of reionisation (for example, Madau et al. 1999). We will therefore verify that analytical reionisation calculations incorporating clumping factors do, in fact, reproduce the behavior of numerical simulations. This check will assess directly the magnitude of the systematic uncertainties inherent in calculations that ignore higher-order effects such as shadowing, source clustering, and the light-travel time of ionising photons.

Having verified that our analytical model reproduces the behavior of our numerical model reasonably well, we will use it to ask how many photons are absorbed by filamentary overdensities in the diffuse IGM. This will quantify the error introduced in models that compute reionisation using IGM density distributions in which the gas is effectively preheated prior to the onset of reionisation.

The outline of our paper is as follows: We begin by introducing our suite of simulations in Section 2. We compare different definitions of the clumping factor in Section 3 and explore how feedback effects modulate their values in Section 4. In Section 5, we incorporate our derived clumping factors into an analytical model for reionisation and test how well analytical models of reionisation reproduce the behavior of numerical calculations. We apply this model to study how many photons are absorbed in filaments in Section 6. Finally, we discuss our results in Section 7 and summarize in Section 8. We introduce two improvements to our treatment of the Eddington tensors in Appendix A and quantify resolution limitations in Appendix B.

## 2 SIMULATIONS

### 2.1 Star Formation and Outflows

Our numerical methods, input physics, and cosmology are similar to Finlator et al. (2011) with the exception of our adopted treatments for an evolving  $f_{\text{esc}}$ , subgrid self-shielding (see below), and two numerical optimizations regarding the treatment of the Eddington tensor field (Appendix A). We model the growth of structure and the feedback processes that couple galaxies with the IGM using our custom version of the cosmological galaxy formation code GADGET-2 (Springel 2005). GADGET-2 implements a formulation of smoothed particle hydrodynamics (SPH) that simultaneously conserves entropy and energy and solves for the gravitational potential with a tree-particle-mesh algorithm. Dense gas cools radiatively using the primordial cooling processes in Table 1 of Katz et al. (1996) and the metal-line cooling tables of Sutherland & Dopita (1993), which assume collisional ionisation equilibrium. In contrast with Katz et al. (1996), however, we evolve the ionisation states of hydrogen and helium and the electron abundance simultaneously with the cooling equations using a nested

subcycling approach whose timestep is limited by a chemical Courant condition; see Finlator et al. (2011) for details. We initialize the IGM temperature and neutral hydrogen fraction to the values appropriate for each simulation's initial redshift as computed by RECFAST (Wong et al. 2008), and we assume that helium is initially completely neutral. We generate the initial density field using an Eisenstein & Hu (1999) power spectrum at redshifts of 249 and 319 for simulations subtending 6 and  $3 h^{-1}$  Mpc, respectively. All simulations assume a cosmology in which  $\Omega_M = 0.28$ ,  $\Omega_\Lambda = 0.72$ ,  $\Omega_b = 0.046$ ,  $h = 0.7$ ,  $\sigma_8 = 0.82$ , and the index of the primordial power spectrum  $n = 0.96$ .

Our goal of modeling galaxies and reionisation simultaneously requires a treatment for the ability of cool gas to form stars. We adopt the subgrid two-phase interstellar medium treatment of Springel & Hernquist (2003), which can be tuned to reproduce the observed relation between the surface densities of gas and star formation (Kennicutt 1998). The physical density threshold for star formation is  $0.13 \text{ cm}^{-3}$ . This value is motivated by observations of a critical density for the onset of star formation and lies within the range at which the thermo-gravitational instability is expected to become active (Schaye 2004). Varying it has only a minor impact on the predicted star formation rate density (Schaye et al. 2010). We account for metal enrichment owing to supernovae of Types II and Ia as well as asymptotic giant branch stars; see Oppenheimer & Davé (2008) for details.

Allowing cold gas to form stars without any feedback invariably results in overproducing the observed reionisation-epoch star formation rate density (Davé et al. 2006; Finlator et al. 2011). The accepted solution is to allow feedback from massive stars to expel star-forming gas from galaxies. As the spatial resolution necessary to form such outflows self-consistently is beyond the reach of current cosmological simulations (Mac Low & Ferrara 1999; Hopkins et al. 2012; Powell et al. 2011), we impose star formation-driven outflows by stochastically applying kicks to star-forming gas particles. The amount of gas kicked per unit stellar mass formed and the kick velocities are adjusted to reproduce the expected scalings from momentum-driven outflows (Murray et al. 2005). We temporarily disable hydrodynamic forces in outflowing gas in order to mimic the way in which outflows escape through holes in higher-resolution simulations (Mac Low & Ferrara 1999). Hydrodynamic forces are restored once the gas has traveled for  $1.95 \times 10^4 \text{ km s}^{-1}/v$  million years (where  $v$  is the kick velocity) or its density drops below 10% of the threshold density for star formation. In practice, this allows outflowing gas to reach galactocentric radii of 50–100 physical kpc and then re-accrete onto the central galaxy on a timescale of  $\sim 1$  Gyr (Oppenheimer & Davé 2008). The distance that outflows travel before re-accreting is large compared to the typical radius at which hydrodynamic forces are enabled, hence the decoupling prescription has little impact on outflows once they escape the ISM. The long re-accretion timescale enables expelled gas to modify the IGM recombination rate and opacity depending on the density of outflowing gas; exploring these possibilities is one of the goals of the present work.

## 4 Gas Clumping at $z \geq 5$

### 2.2 Radiation Transport

We evolve the EUVB on a fixed Cartesian grid using the moments of the cosmological radiative transfer equation. We close the moment hierarchy with a variable Eddington tensor field that is updated periodically via a time-independent ray-casting calculation. This approach accounts accurately for the inhomogeneous opacity field as long as the updates occur frequently enough (Finlator et al. 2009a). Each cell's ionising luminosity is given by a sum over the luminosities of its star-forming gas particles. Their luminosities, in turn, depend on their star formation rates and metallicities, where the metallicity dependence follows an analytical fit to Table 4 of Schaerer (2003).

The fraction of ionising photons that escapes into the IGM in order to participate in reionisation  $f_{\text{esc}}$  is unknown. Recent work suggests that it must vary steeply with redshift in order to bring observations of galaxies, the Lyman- $\alpha$  forest, and the CMB optical depth to Thomson scattering into agreement (Inoue et al. 2006; Finlator et al. 2011; Haardt & Madau 2012; Kuhlen & Faucher-Giguère 2012). We adopt a strongly redshift-dependent  $f_{\text{esc}}$  in our simulations:

$$f_{\text{esc}} = \begin{cases} f_{\text{esc},5} \left(\frac{1+z}{6}\right)^\kappa & z < 10 \\ 1.0 & z \geq 10 \end{cases} \quad (1)$$

Here, the normalization  $f_{\text{esc},5}$  sets the escape fraction at  $z = 5$  and  $\kappa$  controls how strongly  $f_{\text{esc}}$  varies with redshift. For simulations with outflows, we use  $f_{\text{esc},5} = 0.0519$  and  $\kappa = 4.8$ . The normalization is consistent with observations of Lyman break galaxies (Nestor et al. 2011) at  $z \sim 3$ , while the power-law index lies within the range (1–6) that Kuhlen & Faucher-Giguère (2012) indicate is required by observations if the UV luminosity function of galaxies extends to absolute magnitudes of -15 (as required by observations of gamma-ray bursts; Trenti et al. 2012). With this parametrization, the effective optical depth to Lyman- $\alpha$  absorption at  $z = 5$  is  $\tau_\alpha = 3.1$ , marginally consistent with the observed range of 2–3 (Fan et al. 2006). Meanwhile, the predicted ionising emissivities are  $(4.6, 7.0) \times 10^{50} \text{ s}^{-1} \text{ Mpc}^{-3}$  at  $z = (5, 6)$ . At  $z = 5$ , this is consistent with the observed range of  $4.3 \pm 2.6$  (Kuhlen & Faucher-Giguère 2012). At  $z = 6$ , however, it overshoots the observed limit  $< 2.6$ , suggesting that the predicted emissivity should strengthen rather than decline toward the end of reionisation. Note that our model is not unique in failing to reproduce the observed emissivities at  $z = 5$  and  $z = 6$  simultaneously; this evolution is evidently quite strong compared to expectations from models (see, for example, Figure 10 of Kuhlen & Faucher-Giguère 2012 or Figures 8 and 15 of Haardt & Madau 2012). The resulting reionisation history corresponds to an optical depth to Thomson scattering of 0.071, only slightly below the observed  $1\sigma$  confidence intervals of  $\tau_{\text{es}} = 0.088 \pm 0.015$  (Komatsu et al. 2011).

Simulations without outflows require a steeper power-law index and a lower normalization because they overproduce the observed galaxy abundance (Davé et al. 2006). We adopt  $f_{\text{esc},5} = 0.0126$  and  $\kappa = 7.21$ . With this dependence, they predict emissivities of 5.4 and  $7.8 \times 10^{50} \text{ s}^{-1} \text{ Mpc}^{-3}$  at  $z = 5$  and 6, respectively, once again indicating good agreement at  $z = 5$  but incorrect evolution to higher redshift. The predicted Lyman- $\alpha$  optical depth at  $z = 5$  is 2.2, well

within the observed range. The predicted Thomson scattering optical depth is 0.074, within the observed  $1\sigma$  confidence intervals.

The physical interpretation of an evolving  $f_{\text{esc}}$  could take many forms. For example, direct measurements of ionising continuum flux from galaxies at  $z \approx 3$  suggest that current stellar population models underestimate the ionising luminosity of low-metallicity Population II stars (Nestor et al. 2011). If confirmed, this observation could imply an evolutionary trend to higher ionising luminosity at high redshift that mimics an evolving  $f_{\text{esc}}$ . Alternatively, observations indicate that galaxies grow bluer and by inference less dusty beyond  $z = 4$  (Finkelstein et al. 2010); this could drive up  $f_{\text{esc}}$  if dust dominates the absorbing column (although Gnedin et al. 2008 argue that  $f_{\text{esc}}$  is dominated by the ISM's geometry rather than its dust content). Finally, galactic outflows can fill a halo with gas, potentially creating a significant optical depth to ionising photons. If these screens grow weaker with increasing lookback time, then the fraction of ionising photons that travel unimpeded from the edge of a galaxy's ISM to the virial radius could grow accordingly; this would mimic an evolving  $f_{\text{esc}}$ . We defer detailed investigation of these possibilities to future work; for the present, we simply choose a parametrization that will improve the realism of our simulated density field. We discuss drawbacks to this approach in Section 7.

We assume that each photoionisation deposits 4.08 eV of latent heat into the IGM. This heats newly-reionised gas to  $\approx 15,000$  K, which lies within the range that is expected for a Population II stellar spectrum. Each radiation transport cell's opacity is given by its volume-averaged neutral hydrogen density (counting only the gas that is not self-shielded; see below) multiplied by the cross-section for ionisation at an energy of 17.68 eV.

### 2.3 Self-shielding Within Overdense Systems

Two central challenges in using numerical models to compute the IGM recombination rate involve deciding which gas belongs to the IGM and ensuring that overdense gas is allowed to self-shield against the EUVB.

The need to define the IGM in a physically-motivated way is illustrated by the results of the pathbreaking radiation hydrodynamic simulations of Gnedin & Ostriker (1997) and Gnedin (2000). These simulations faithfully modelled the growth of ionised regions as well as the hydrodynamic response of the heated gas. However, the volume-averaged recombination rates were computed over all the gas in the simulation, leading to recombination rates that were much larger than what is plausible for the moderately overdense  $(\rho/\langle\rho\rangle < 100)$  gas that is conventionally associated with the IGM. It also led to a strong dependence of the recombination rate (and the reionisation photon budget) on the simulation dynamic range, with higher-resolution simulations absorbing more photons (Gnedin 2000).

Improvement results from using simple density cuts to isolate low-density gas in post-processing from precomputed simulations. This approach has been applied both to N-body simulations (Iliev et al. 2007; Raičević & Theuns 2011) and to cosmological hydrodynamic simulations (Pawlik et al. 2009; Shull et al. 2012). It reduces  $C$  as expected, but the results remain dependent on the precise cuts used to de-

fine the IGM. In the case of hydrodynamic simulations, additional uncertainty results from assuming a homogeneous EUVB, which may pressurize regions that are in reality dense enough to self-shield and does not account for the spatial inhomogeneity of reionisation. Both of these effects could modify the IGM's density and temperature fields (Furlanetto & Oh 2009), which in turn impact the recombination rate.

In this work, we combine these two approaches by using radiation hydrodynamic simulations of the reionisation epoch that incorporate a physically-motivated treatment for self-shielding within overdense regions. Effectively, we incorporate directly into our simulations the definition that *the IGM consists only of that gas whose ionisation state is not determined by local sources*; denser systems are allowed to remain neutral. By isolating overdense systems from the EUVB, we (1) improve the accuracy of our predicted density and temperature fields; (2) prevent our simulations from overestimating the reionisation photon budget; and (3) obviate the need to choose a self-shielding threshold overdensity (such as 50 or 100) when deriving the IGM recombination rate in post-processing: the threshold evolves on-the-fly in a way that follows the (local) gas temperature and (global) EUVB amplitude.

We model the self-shielding of overdense regions using a subgrid recipe because our radiation transport solver does not resolve them spatially even though our hydrodynamic solver does. For example, Lyman limit systems (with neutral column densities of  $\sim 10^{17} \text{ cm}^{-2}$ ) have a characteristic size of 10 physical kpc (Schaye 2001; McQuinn et al. 2011) and must be treated with 5–10 resolution elements (Gnedin & Fan 2006). In our higher-resolution simulation (r6n256wWwRT32), the gravitational softening length is 0.1 kpc (Plummer equivalent) while the radiation transport solver achieves a resolution of 38 kpc (where both numbers are in physical units at  $z = 6$ ). Hence while our hydrodynamic solver certainly resolves Lyman limit systems, our radiation solver's resolution is a factor of  $\approx 20$  too coarse.

We overcome this limitation by following the argument of Schaye (2001). This work showed that, if Lyman limit systems are in hydrostatic equilibrium, then they can be identified directly with gas above a critical overdensity (see also McQuinn et al. 2011). The threshold,  $\Delta_{\text{lls}}$ , is given by:

$$\Delta_{\text{lls}} = 25 \left( \frac{T}{10^4 \text{ K}} \right)^{0.3} \left( \frac{1+z}{7} \right)^{-3} \left( \frac{\Gamma_{\text{HI}}}{2 \times 10^{-13}} \right)^{\frac{2}{3}} \quad (2)$$

at the Lyman limit. Here we have assumed that hydrogen is fully ionised and helium is singly-ionised, as expected for soft ionising sources (Ciardi et al. 2012; Friedrich et al. 2012).  $\Gamma_{\text{HI}}$  represents the volume-averaged EUVB and does not include the influence of local sources. This estimate is within a factor of two of what others have found (for example, Bolton & Haehnelt 2007). We do not expect slight differences owing, for example, to our assumed recombination rate to affect our results. We enforce  $\Delta_{\text{lls}} \geq 2$  in order to preserve self-shielding at  $z \gg 10$ ; this floor is not expected to affect our results. We use this approximation to model subgrid self-shielding as follows: For each gas particle whose density lies below the physical threshold for star formation ( $0.13 \text{ cm}^{-3}$ ), we compute  $\Delta_{\text{lls}}$  using the local gas temperature, the case-A recombination rates, and the global mean

ionisation rate per hydrogen  $\Gamma_{\text{HI}}$ . We then assume that the optical depth to the diffuse background varies as a power-law of the overdensity

$$\tau_T = \begin{cases} 0 & \Delta < \Delta_{\text{lls}} \\ \tau_0 \left( \frac{\Delta}{\Delta_{\text{lls}}} \right)^b & \Delta \geq \Delta_{\text{lls}} \end{cases} \quad (3)$$

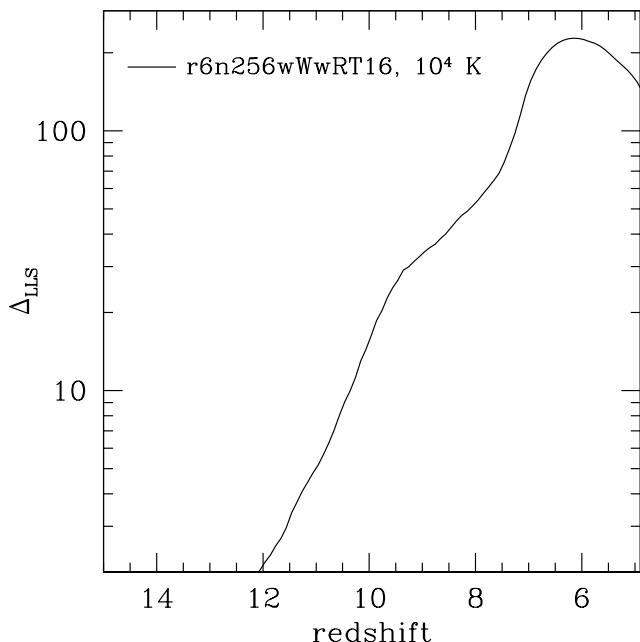
In the optically-thin limit,  $\tau_0 = 1$  and  $b = 1.5$  (Schaye 2001). In reality,  $\tau_T$  increases more rapidly with density because the neutral fraction in the partially shielded outskirts of an absorber is larger than in the optically-thin case. We have experimented with both  $b = 1.5$  and  $b = 3$  and found little difference in practice. We adopt the fiducial choices  $\tau_0 = 1.0$  and  $b = 3$  in this work.

We implement Equations 2–3 directly into our simulations so that overdense gas is shielded from the EUVB on-the-fly. Our use of different discretizations for the radiation and hydrodynamic solvers requires that we do so in two ways. When updating the ionisation state of a given particle, we attenuate the radiation field of its host cell by the factor  $e^{-\tau_T}$ , where  $\tau_T$  is evaluated using the global EUVB and the particle's temperature. Conversely, when assembling the opacity field on the radiation solver's grid, we reduce each particle's contribution to its host cell's opacity by the same factor. We include the opacity of each gas particle's self-shielded region by treating it as an opaque sphere with volume equal to  $f x_{\text{HI}} (1 - e^{-\tau_T}) \Delta V$ , where  $\Delta V$  is the ratio of the gas particle's mass to its density,  $x_{\text{HI}}$  is its neutral fraction, and  $f$  is a parameter that is tuned via high-resolution calculations to 1/8. For example, if  $e^{-\tau_T} = 0.5$  for a particle, then its ionisation rate is 50% of that of its host cell's, and only 50% of its neutral hydrogens contribute to the cell's opacity. We also augment the cell's opacity by the amount  $4.84(0.5x_{\text{HI}}\Delta V)^{-2/3}/A_{\text{cell}}$ , where  $A_{\text{cell}}$  is the cell's area. Note that the opacity from partially-neutral gas particles whose overdensity falls below the self-shielding threshold is included directly in the radiation transport solver as in our previous work.

Intuitively, this treatment divides self-shielded regions into an optically thin skin and an optically-thick core. The overdensity range over which a region transitions from optically-thin to optically-thick follows from Equation 3. The opacity owing to the optically-thin region is distributed uniformly throughout the radiation transport cell while the core is treated as a photon sink. This approximation preserves the ability of dense gas to remain self-shielded within a coarse radiative transfer grid. It will enable radiative hydrodynamic simulations to subtend cosmological volumes even as they treat the radiation field's small scale structure with sufficient detail to model observables such as the abundance of low-ionisation metal absorbers and the post-reionisation IGM opacity.

In reality, regions that are dense enough for star formation to occur ( $\Delta > 1000$ ) contain a mixture of ionised and neutral gas owing to the local radiation field from massive stars. They remain entirely neutral in our simulations owing to our self-shielding prescription. Accounting for their recombinations would lead to much larger clumping factors (Gnedin & Ostriker 1997; Gnedin 2000). However, the corresponding absorptions should be identified with an ionising escape fraction  $f_{\text{esc}}$  that is less than unity rather than an increased clumping factor because they are not domi-

## 6 Gas Clumping at $z \geq 5$



**Figure 2.** The threshold overdensity for self-shielding  $\Delta_{lls}$  as a function of redshift within our fiducial simulation at a characteristic temperature of  $10^4$ K.  $\Delta_{lls}$  peaks at over 200 during the heyday of reionisation before shrinking to roughly 150 by  $z = 5$  owing to declining escape fractions.

nated by the EUVB. We encapsulate their impact via an assumed  $f_{\text{esc}} \leq 1$  that is tuned using observational insight (see above) and defer the challenging task of modeling  $f_{\text{esc}}$  directly to future work.

Figure 2 shows how  $\Delta_{lls}$  evolves in our fiducial simulation for a characteristic temperature of  $10^4$ K. Note that this figure is illustrative; in reality our simulations compute this threshold on-the-fly using the local temperature.  $\Delta_{lls}$  grows from values near unity at the onset of reionisation to values near 200 by the time that reionisation has completed. Gas that is denser than this threshold sees an attenuated EUVB, suppressing both the reionisation photon budget and the volume-averaged recombination rate.

In detail, the values in Figure 2 may be somewhat overestimated (see, for example, McQuinn et al. 2011). This occurs because our monochromatic radiation transport solver uses a smaller photoionisation cross section  $\sigma$  than the value at the Lyman limit  $\sigma_{LL}$  in order to match the assumed photoheating rate of 4.08 eV per photoionisation. In particular,  $\Delta_{lls}$  varies with the photoionisation cross section  $\sigma$  as  $(\sigma/\sigma_{LL})^{-2/3}$  (Schaye 2001), hence our simulations may overestimate  $\Delta_{lls}$  by a factor of  $\approx 1.7$ . This is roughly the difference between our value at  $z = 6$  and that of McQuinn et al. (2011). The overall impact of this limitation on our inferred value of  $C$  is weak; we will quantify it in Section 7.

Table 1 introduces our simulation suite. The simulation names encode the physics and numerical resolutions. For example, the r6n256wWwRT16 simulates a  $6h^{-1}\text{Mpc}$  volume (r6) with  $2 \times 256^3$  particles (n256), includes outflows (wW) and evolves the EUVB on a Cartesian grid with  $16^3$  cells (wRT16). The first four simulations explore the sensitivity of gas clumping to photoheating and outflows. The

next three explore convergence with respect to our radiation and gas solvers as well as the cosmological volume. The final simulation assumes an optically-thin EUVB for comparison with previous work. Throughout this work, we will use the r6n256wWwRT16 simulation as our fiducial case.

Our simulations allow us to explore how the IGM's temperature impacts its recombination rate within the context of a model that treats photoheating from Population II stars self-consistently (Section 4.1). As a check on whether our predictions are realistic, we compare the temperature at mean density to recent observations. Bolton et al. (2012) used measurements of the Doppler widths of Ly $\alpha$  absorption lines along the line of sight to seven quasars to infer that the IGM temperature at mean density is  $\log(T) = 3.85 \pm 0.08$  at  $z = 6.08 \pm 0.33$  (their Table 3, fiducial model). Our fiducial simulation predicts  $\log(T, z = 6) = 3.93 \pm 0.10$ . Here, we report the median  $\log(T)$  over particles with overdensity  $-0.2 < \log(\Delta) < 0.2$ ; switching from the median to the mean changes results by  $\sim 0.1\%$  because the scatter in the temperature distribution is not large. This comparison is of course incomplete because it does not account for a variety of observational systematics. Nonetheless, the fact that the simulation reproduces the observed value within the errors indicates that the predicted temperature fields are plausible.

## 3 WHAT IS GAS CLUMPING?

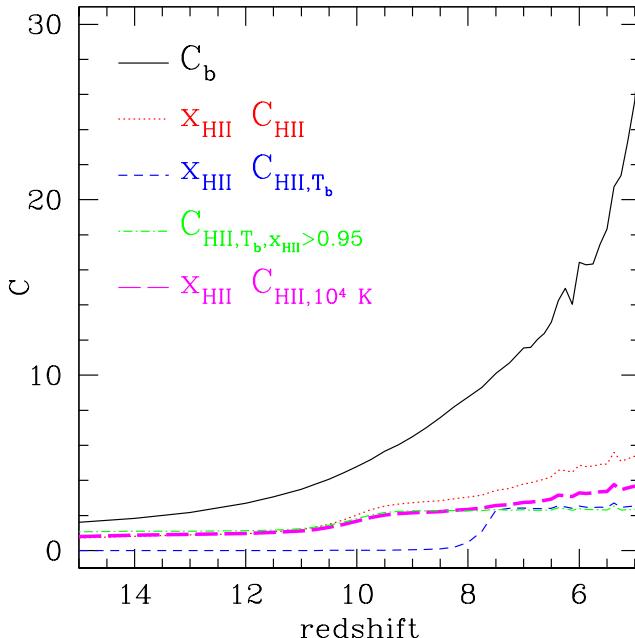
The motivation for defining clumping factors comes from a desire to take the volume-weighted mean of the ionisation rate equations. For example, the rate of change of the neutral hydrogen abundance  $n_{\text{H I}}$  is given by

$$\frac{dn_{\text{H I}}}{dt} = \Gamma_{\text{HI}} n_{\text{H I}} + k_2 n_{\text{H II}} n_{\text{e}} - k_1 n_{\text{H I}} n_{\text{e}}, \quad (4)$$

where  $\Gamma_{\text{HI}}$  is the photoionisation rate per hydrogen atom,  $k_1$  is the collisional ionisation rate, and  $k_2$  is the (case A) recombination rate. Each term on the right side is nonlinear, hence, in principle, it is not possible to compute their spatial averages unless the cross-correlation of the abundances of  $n_{\text{H I}}$  and  $n_{\text{H II}}$  with each other as well as with the radiation and temperature fields are known (since  $k_1$  and  $k_2$  are temperature-dependent). It is convenient to encapsulate this subgrid information with clumping factors. For example, the spatially averaged recombination rate is approximated as

$$\langle k_2 n_{\text{H II}} n_{\text{e}} \rangle \equiv C_{\text{H II}, \text{T}_b} \langle n_{\text{H II}} \rangle \langle n_{\text{e}} \rangle \langle k_2 \rangle, \quad (5)$$

where angle-brackets indicate averages over the entire simulation volume. Similar clumping factors could be defined for each of the terms on the right-hand side of Equation 4. In practice, however, analytical calculations bypass the need for the first term's clumping factor by assuming that the volume-averaged  $\Gamma_{\text{HI}}$  equals the ionising emissivity divided by the mean hydrogen abundance. Meanwhile, the third term can be neglected because  $\Gamma_{\text{HI}} \gg k_1 n_{\text{e}}$  except within shocks. This means that, along with the ionising escape fraction,  $C_{\text{H II}, \text{T}_b}$  is one of the major uncertainties that hamper efforts to connect the observed abundance of ionising sources (such as Lyman Break galaxies) with the ionisation state of the IGM (see Kohler et al. 2007 for an expanded discussion of the other clumping factors). In Section 5, we will show



**Figure 3.** The evolution of four different clumping factors in our fiducial simulation volume. The clumping factors of all gas whose density exceeds the star formation threshold reaches 15 by  $z = 6$ . Accounting for the fact that overdense gas is self-shielded and that ionised gas is somewhat warmer than  $10^4$  K leads to much lower clumping factors (see text for discussion).

that the numerical results may be reproduced by an analytical model that considers only the recombination clumping factor even though the full numerical calculation accounts for all three terms.

Equation 5 shows that the clumping factor depends on the density, ionisation, and temperature fields. Partial clumping factors are often used when one or more of these fields is unavailable. In Figure 3, we show how the choice of definition impacts the clumping factor’s value in our fiducial simulation. The black solid curve shows the evolution of the clumping factor averaged over all gas whose density lies below the adopted threshold for star formation ( $0.13 \text{ cm}^{-3}$ ),  $C_b \equiv \langle \rho^2 \rangle / \langle \rho \rangle^2$ . It increases monotonically as overdense regions collapse, with a brief enhancement around the reionisation redshift. The enhancement occurs when the EUVB grows strong enough to photoevaporate gas from photosensitive halos (Finlator et al. 2011), temporarily boosting the mass of gas whose density lies just below the threshold for star formation. The clumping factor of all gas (including star forming gas) decreases at this epoch as expected, but the clumping factor of gas that lies outside of galaxies—that is,  $C_b$ —increases until the newly-pressurized gas has expanded into hydrostatic equilibrium.

The clumping factor within ionised regions—which is the important quantity for reionisation—is expected to be different from  $C_b$  because reionisation begins in overdense regions and progresses into voids (Furlanetto et al. 2004; McQuinn et al. 2007) in such a way that the last regions to remain neutral are overdense (Miralda-Escudé et al. 2000). Furthermore, gas that is self-shielded does not contribute to recombinations. To explore these factors, we compute the

clumping factor of electrons and protons

$$C_{\text{HII}} \equiv \langle n_e n_{\text{HII}} \rangle / \langle n_e \rangle \langle n_{\text{HII}} \rangle. \quad (6)$$

$C_{\text{HII}}$  can be computed in either of two ways. One way is to compute the averages over all ionised gas. Unfortunately, this introduces uncertainty owing to the choice of ionisation threshold. For example, averaging only over gas particles that are more than 50% ionised gives a slightly different answer than averaging only over particles that are more than 95% ionised because the simulated IGM is not a strictly two-phase medium. The alternative is to take averages over the entire simulation volume and then weight by the ionised volume fraction  $x_{\text{HII},v}$ . This can be understood by recalling that, in a partially-ionised universe with a uniform density,  $C_{\text{HII}} = 1/x_{\text{HII},v}$ . Therefore,  $x_{\text{HII},v} C_{\text{HII}}$  indicates the recombination rate in units of the recombination rate of a homogeneous universe that has the same ionisation state, which is the goal of the clumping factor.

We show this quantity with a red dotted line. The tendency for overdense regions to remain neutral suppresses  $x_{\text{HII},v} C_{\text{HII}}$  below  $C_b$  at all times, with reionisation-epoch values that lie below  $\approx 6$ . This difference immediately reveals the importance of self-shielding, which creates a boundary between the ionised IGM and the neutral or locally-ionised regions near halos. It can be modeled directly by spatially resolving physical scales of  $\sim 1$  kpc (Gnedin & Fan 2006). However, this incurs significant computational expense, limiting the simulation’s cosmological volume. Our approach of coupling a somewhat coarse grid for the radiation solver with a subgrid prescription for self-shielding may open up the possibility of simulating cosmological reionisation within large volumes while treating the IGM’s thermal and ionisation states faithfully.

$C_{\text{HII}}$  is an improved description of the volume-averaged recombination rate over  $C_b$ , but it remains incomplete because the recombination rate  $k_2$  also depends on the gas temperature. To illustrate this, we use a blue dashed curve to show the temperature-corrected clumping factor  $x_{\text{HII},v} C_{\text{HII},T_b}$ , computed following Equation 5 (referred to by Shull et al. 2012 as  $C_{\text{RR}}$ ). This curve remains near zero until reionisation is well under way because the recombination rate per proton is spatially anti-correlated with the abundance of ionised gas. Once  $x_{\text{HII},v}$  exceeds  $\approx 0.5$ , it begins to climb because  $\langle k_2 \rangle$  is no longer suppressed by cold neutral regions. However, it does not reach  $C_{\text{HII}}$  owing to the lingering presence of self-shielded regions. Following reionisation,  $x_{\text{HII},v} C_{\text{HII},T_b}$  climbs slowly because the volume-weighted mean temperature is supported by the slow photo-heating of filaments, which suppresses  $\langle k_2 \rangle$ . It never exceeds 3, reflecting the tendency of photoheating to suppress recombination rates.

$x_{\text{HII},v} C_{\text{HII},T_b}$  provides a reasonable description of the clumping factor within ionised regions once reionisation is well under way, but it is less informative at early times ( $x_{\text{HII},v} < 0.1$ ) when the high recombination rate in the predominantly cold, neutral IGM suppresses  $C_{\text{HII},T_b}$  despite the high clumping factor of ionised regions. We may compute a more informative clumping factor by re-evaluating  $C_{\text{HII},T_b}$  using only those regions (that is, SPH particles) whose ionised mass fraction is greater than 0.95. This clumping factor (green dot-dashed) is somewhat larger at earlier times because it is not suppressed by the high recombination

## 8 Gas Clumping at $z \geq 5$

rates per proton in regions that have not yet been heated. It merges with  $x_{\text{H II},v} C_{\text{H II},T_b}$  once  $x_{\text{H II},v} \sim 1$ . A drawback of  $C_{\text{H II},T_b}$  is that its value depends on the threshold ionisation fraction that is used to compute it.

While  $C_{\text{H II},T_b}$  emerges naturally from a spatial average of Equation 4, it is difficult to compare with observations because the mean IGM temperature (and hence  $\langle k_2 \rangle$ ) at  $z = 6$  remains poorly-constrained; this is of course equally true even if  $C_{\text{H II},T_b}$  is computed only over regions whose ionisation fraction exceeds some threshold. To resolve these problems, we also compute the ‘‘observational temperature-corrected clumping factor of ionised gas’’

$$C_{\text{H II},10^4\text{K}} \equiv \frac{\langle n_e n_{\text{H II}} k_2(T) \rangle}{\langle n_e \rangle \langle n_{\text{H II}} \rangle k_2(10^4\text{K})} \quad (7)$$

in which we replace the volume-averaged recombination rate  $\langle k_2 \rangle$  in the denominator with the recombination rate at  $10^4\text{K}$ . This clumping factor (referred to by Gnedin 2000 as  $C_{\text{HII}}$ ) expresses the mean recombination rate without requiring knowledge of the topology of reionisation or the temperature of ionised gas. It is slightly higher than  $C_{\text{H II},T_b}$  following reionisation, reflecting the tendency for gas to cool below  $10^4\text{K}$  once reionisation is complete (that is,  $\langle k_2 \rangle > k_2(10^4\text{K})$  for  $z < 6$ ). The function

$$x_{\text{H II},v} C_{\text{H II},10^4\text{K}} = 9.25 - 7.21 \log_{10}(1+z) \quad (8)$$

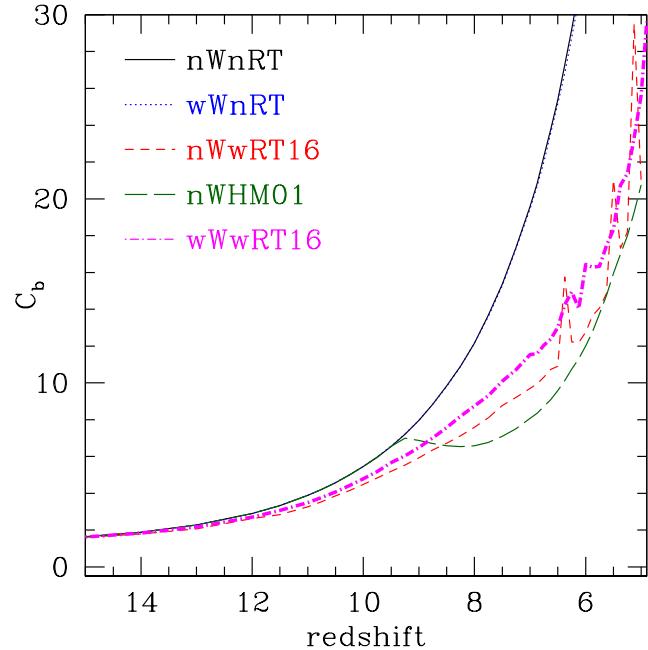
fits its evolution to within 30% for  $z \leq 15$  and 10% for  $z \leq 10$ .

Note that uncertainties in the recombination rate translate directly into uncertainties in the clumping factor. For example, our adopted (case A) recombination rate is  $\sim 20\%$  higher than that of Hui & Gnedin (1997) in the temperature range  $10^3\text{--}10^5\text{K}$ . This means that ionisation fronts do not penetrate as far into overdense regions in our simulation as they would if we assumed the Hui & Gnedin (1997) rates, reducing the inferred clumping factor of ionised gas by the same ratio. The resulting uncertainty does not affect our results and is trivial to adjust for: Simply divide our inferred clumping factors by the ratio of the alternative recombination coefficient to our adopted coefficient at  $10^4\text{K}$ .

In summary, Figure 3 indicates that the clumping factor’s value depends rather strongly on its definition. Taking self-shielding into account suppresses it from  $\approx 30$  to  $\approx 6$  at  $z = 5$  because the most overdense regions are self-shielded and do not contribute to the IGM recombination rate. Taking the IGM’s temperature into account further suppresses clumping to 2–4 because photoionised regions are generally warmer than  $10^4\text{K}$ .

## 4 THE ROLE OF FEEDBACK

Galactic outflows and photoheating both affect gas clumping. Heating suppresses gas clumping by smoothing density fluctuations on scales smaller than the local Jeans length (Efstathiou 1992). Outflows increase IGM clumping partly by boosting the amount of dense gas that resides within halos but outside of galaxies and partly (for a given  $f_{\text{esc}}(z)$ ) by delaying reionisation and heating. The latter effect is a straightforward consequence of suppressing star formation. The former is expected based on our



**Figure 4.** The impact of photoheating and outflows on the overall IGM baryonic clumping factor  $C_b$ . All simulations subtend a  $6h^{-1}\text{Mpc}$  volume discretized with  $2 \times 256^3$  particles (r6n256); see Table 1. Outflows have negligible impact on  $C_b$  in the absence of an EUVB, but when acting in concert with an EUVB their impact is stronger than uncertainties owing to limitations in our particular treatment of the EUVB.

previous finding that outflows boost the number of ionising photons required to achieve reionisation (Finlator et al. 2011) although it may in reality be weaker if ejected gas self-shields. In this section, we quantify how these processes modulate gas clumping.

### 4.1 Averaging over the IGM

We begin by discussing the separate impacts of outflows and photoheating on the IGM clumping factor. In order to facilitate this discussion, we revert to the definition  $C_b$ , which uses all gas that is too diffuse to form stars (Section 2) without reference to its ionisation state or temperature. While this definition includes gas that is dense enough to self-shield and is therefore an imperfect estimate of the IGM recombination rate, it allows us to compare simulations with and without an EUVB.

We use solid black and dotted blue curves in Figure 4 to show how  $C_b$  evolves in simulations without and with outflows in the absence of an EUVB. They are nearly coincident, indicating that outflows affect the gas density distribution only weakly if they do not couple to an EUVB. Adopting a spatially-homogeneous Haardt & Madau (2001, hereafter HM01) EUVB dramatically suppresses the clumping factor following  $z = 9$  (long-dashed green; see also Pawlik et al. 2009).

Our self-consistent simulation omitting outflows (r6n256nWwRT16, short-dashed red) predicts a reionisation history in which the neutral hydrogen fraction drops to 50% at  $z = 9.2$ , quite similar to the redshift at which the

HM01 EUVB reionises the Universe ( $z = 9$ ). Prior to  $z = 5$ , its  $C_b$  exceeds the HM01 case owing to its more extended reionisation history. By  $z = 5$ , its  $C_b$  converges to the HM01 case although it shows several spikes that are likely associated with photoevaporation of star-forming gas from low-mass ( $M_h < 10^9 M_\odot$ ) haloes in regions that are just being reionised. This suggests that  $C_b$  is weakly sensitive to the reionisation history at early times but converges to expectations from calculations using a spatially-uniform EUVB fairly rapidly (probably on the sound-crossing time of the filaments).

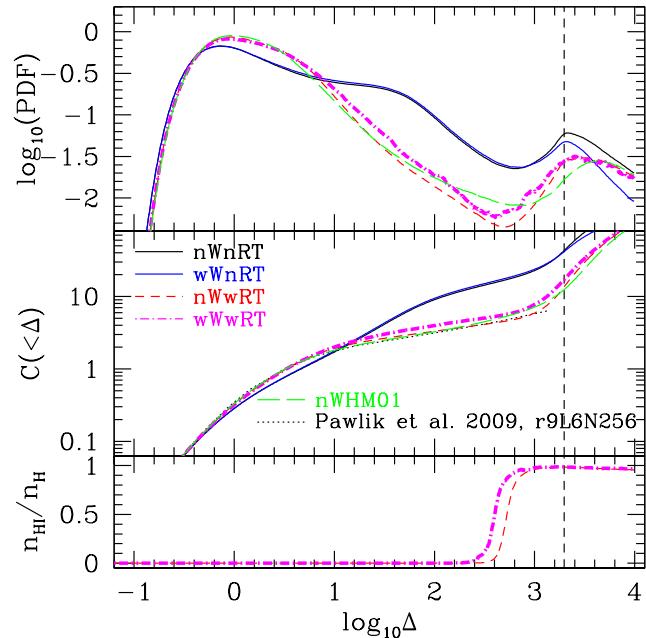
Turning to our fiducial simulation (magenta dot-dashed), which does include outflows, we find that its  $C_b$  is higher than the no-wind case at all redshifts. This model has significantly less star formation than the no-wind model, but it also assumes a higher ionising escape fraction in order to compensate (Section 2.2). Combining these factors leads to a reionisation history in which the neutral fraction drops below 50% at  $z = 8.9$ , only slightly later than the other models. Despite their similar reionisation histories, the wind model's  $C_b$  is roughly 30% higher at  $z = 5$ . Hence although the reionisation histories of the wind and no-wind models are not exactly the same, Figure 4 confirms that outflows enhance  $C_b$  by boosting the amount of gas that is overdense but not star-forming.

We draw several conclusions from Figure 4. First, outflows do not affect  $C_b$  unless they are heated by an EUVB. Intuitively, outflows consist largely of material entrained from their host galaxies' interstellar media, hence they remain relatively cold and dense unless they are heated by an EUVB. Second, an EUVB heats and evaporates gas out of shallow potential wells such as filaments and minihalos, suppressing  $C_b$ . Finally, a better understanding of the nature of galactic outflows is *at least* as important to our understanding of the IGM density structure as the nature of the EUVB. In Section 4.2, we will ask whether this uncertainty impacts the recombination rate of diffuse gas.

## 4.2 The recombination rate as a function of density

Figure 4 suggests that outflows and photoheating modulate the IGM's density structure, but their impact on the IGM recombination rate remains unclear because  $C_b$  averages over gas that is in reality self-shielded. In order to understand how feedback processes modulate the abundance of gas at different densities, we show the mass-weighted probability density function of gas (PDF) at  $z = 6$  as a function of overdensity  $\Delta \equiv \rho/\langle\rho\rangle$  in the top panel of Figure 5. Comparing the blue and black solid curves to the lower curves shows how photoheating smooths over filaments and removes gas from low-mass haloes (see also Pawlik et al. 2009). At densities higher than the adopted threshold for star formation (vertical dashed line), the PDF is depleted by star formation and outflows.

Above  $\Delta \sim 240$ , both reionisation simulations (short-dashed red and dot-dashed magenta) show an increasing gas abundance because the EUVB is attenuated in these regions owing to self-shielding. Outflows modestly increase the gas abundance throughout the region  $10 < \Delta < 1000$ . The long-dashed green curve shows the PDF from the r6n256nWHM01 simulation, which omits outflows



**Figure 5.** (top) The gas density PDF in units of mass-weighted probability per unit  $\log_{10}(\Delta)$  for six different simulations as indicated by the legend in the middle panel at  $z = 6$ ; all simulations subtend volumes with side length  $6h^{-1}\text{Mpc}$  and resolve halos of total mass  $1.4 \times 10^8 M_\odot$ . The vertical dashed line indicates the threshold density for star and wind formation ( $0.13\text{cm}^{-3}$ ). Photoionisation heating pressurizes gas with densities  $\Delta > 10$ , and outflows transfer gas from galaxies into the IGM. (middle) The cumulative baryonic clumping factor as a function of maximum overdensity compared to the fiducial simulation of Pawlik et al. (2009) (black dotted). An EUVB suppresses the clumping of all gas, with minor differences owing to the details of outflows and reionisation. (bottom) The neutral hydrogen fraction as a function of overdensity. Gas above a threshold overdensity (which is related to  $\Delta_{\text{lhs}}$ ) is neutral owing to our self-shielding model.

and assumes a spatially-homogeneous EUVB (with no self-shielding). Its gas abundance is suppressed with respect to the self-shielding calculations in regions where  $\Delta > 1000$  and enhanced near  $\Delta \sim 200$  because the EUVB evaporates dense gas that should in reality be shelf-shielded.

Differences in the gas density PDFs translate into different amounts of gas clumping. In the middle panel of Figure 5, we show the cumulative clumping factor as a function of the threshold overdensity  $C(< \Delta_{\text{thr}})$ , defined following Pawlik et al. (2009):

$$C(< \Delta_{\text{thr}}) \equiv \frac{\int_0^{\Delta_{\text{thr}}} d\Delta \Delta^2 \mathcal{P}(\Delta)}{\int_0^{\Delta_{\text{thr}}} d\Delta \mathcal{P}(\Delta)} \quad (9)$$

Here,  $\mathcal{P}(\Delta)$  is the volume-weighted probability distribution function of gas density. Note that each curve's value at the star formation threshold equals the corresponding simulation's  $C_b$  at  $z = 6$  in Figure 4. The solid blue and black curves confirm that outflows have little impact on the clumping factor without an EUVB. Adopting a spatially-homogeneous EUVB (green long-dashed) reproduces the Pawlik et al. (2009) result (black dotted) as expected, with slight differences likely owing to their addition of extra heating around the reionisation redshift. Our

radiation-hydrodynamic simulations again predict very similar clumping factors. Outflows boost clumping as expected from the top panel, but the change is slight. In particular, we confirm the finding by Pawlik et al. (2009) that outflows change  $C(< \Delta_{\text{thr}})$  by  $< 10\%$  when averaged over gas with  $\Delta < 100$ . This is somewhat nontrivial given that our outflow prescription ejects roughly twice as much gas per unit stellar mass formed as theirs at this redshift. Overall, the impact of changing the EUVB and the outflow model on  $C(< \Delta_{\text{thr}})$  is quite modest. The fact that the gas density PDF at low densities ( $\Delta < 200$ ) is relatively robust to these factors indicates that the recombination rate in the diffuse IGM—which is the important quantity for understanding the progress of reionisation—is only weakly sensitive to the uncertain details of star formation and feedback. At higher densities, the impact of feedback and self-shielding is much stronger (top panel), but recombinations in these regions do not count toward the IGM recombination rate as their ionisation state is dominated by local sources.

The ability of outflows to boost the abundance of neutral gas at moderate overdensities ( $\Delta > 240$ ) rather than only at the outskirts of galaxies ( $\Delta > 1000$ ) suggests that they could function as an additional population of Lyman limit systems residing near or within halos. If so, then they would have to be self-shielded against the EUVB. We show in the bottom panel the neutral hydrogen fraction as a function of  $\Delta$  for our fiducial simulation. The neutral fraction jumps from 0 to 1 at  $\Delta \approx 240$ , indicating that outflows could indeed be partially self-shielded. This possibility is consistent with direct observations of cold gas in galactic outflows (for example, Martin 2005; Weiner et al. 2009) as well as recent theoretical results (van de Voort et al. 2012).

In summary, an EUVB significantly suppresses IGM clumping by pressurizing moderately overdense regions while outflows modestly boost the recombination rate by returning gas from galaxies into the IGM. Meanwhile, self-shielding hides the most overdense gas from the EUVB. This boosts the mass fraction of neutral gas while suppressing the clumping factor in ionised regions. Outflows are at least partially self-shielded, suggesting that they could contribute to the IGM opacity.

## 5 APPLICATION I: PHOTON-COUNTING

Analytical models of reionisation use clumping factors to account for IGM inhomogeneities (Madau et al. 1999), prompting the need for theoretical insight into the nature of clumping. We have discussed the various effects that modulate the value of the clumping factor that we derive from our simulations. In order to close the loop, we now ask whether these clumping factors do indeed constitute a good model for the IGM recombination rate during reionisation. We do this by using an analytical model to explore how well several different clumping factors fare in reproducing the reionisation history of our r6n256wWwRT16 simulation. We will show that clumping factors perform reasonably well, and that they fare better if they incorporate information from the simulated temperature and ionisation fields.

We use as our model Equation 20 of Madau et al. (1999) (this can also be derived by taking the spatial average of Equation 4 and normalizing by the mean hydrogen number

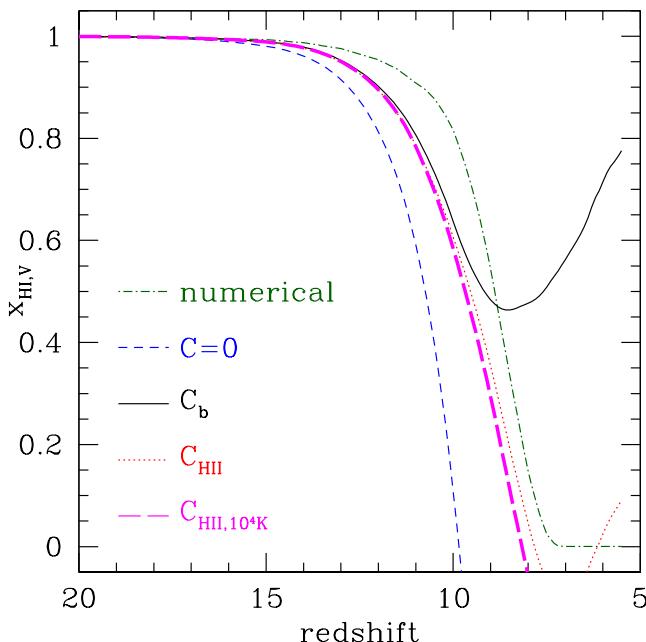
density):

$$\frac{dx_{\text{HI},\text{V}}}{dt} = -\dot{n}_\gamma + \text{recombinations.} \quad (10)$$

Here,  $x_{\text{HI},\text{V}} \equiv \langle n_{\text{HI}} / n_{\text{H}} \rangle$  is the volume-averaged neutral hydrogen fraction while  $\dot{n}_\gamma$  is the ionising luminosity per hydrogen atom into the IGM. We will use the value for  $\dot{n}_\gamma$  that is predicted directly by our simulation. The form of the recombination term varies depending on the clumping factor definition.

There are three caveats to this widely-used formalism. The first is that the second term represents a *mass-average* while the third term is generally computed in a *volume-averaged* way (see, for example, Equation 5). To see this, recall that the second term is meant to model the rate of growth of the ionised volume fraction (Madau et al. 1999). The rate at which an ionised volume  $V_I$  grows depends on the ratio of the ionising luminosity  $\dot{N}_\gamma$  to the gas density at the position of the ionisation front  $n_H$ ,  $\dot{V}_I = \dot{N}_\gamma / n_H$ . In practice, however, Madau et al. (1999) substituted the volume-averaged gas density  $\langle n_H \rangle$  into the denominator. This approximation is only accurate if the density is homogeneous or if the sources are widely separated (as in the case of helium reionisation; McQuinn et al. 2009), hence it is equivalent to assuming that the neutral mass and volume fractions are equal,  $x_{\text{HI},\text{M}} = x_{\text{HI},\text{V}}$ . It is possible to use knowledge of the distribution of gas densities and the topology of reionisation (for example, Miralda-Escudé et al. 2000) to compute the  $\dot{V}_I$  from the ionising emissivity (or, alternatively, to compute the recombination rate from the ionised mass fraction), but these steps are omitted in analytical calculations. Owing to this approximation, we will refer only to the neutral fraction  $x_{\text{HI}}$  rather than to  $x_{\text{HI},\text{M}}$  or  $x_{\text{HI},\text{V}}$  within the context of our analytical model. The second caveat is that Equation 10 omits collisional ionisations because they are nontrivial to model analytically. Finally, the third term requires knowledge of the IGM’s density and temperature distributions that is generally encapsulated via clumping factors. All of these simplifications are relaxed in our numerical calculations. In this section, we will use our simulations to ask whether they are indeed valid.

We compare the resulting reionisation histories with our numerical simulation in Figure 6. As a baseline, the short-dashed blue curve illustrates the reionisation history if there are no recombinations ( $C = 0$ ). Comparing it with the dot-dashed green curve indicates that recombinations delay the completion of reionisation from  $z \approx 10$  to  $z \approx 7.5$ . The delay is expected given that this simulation requires 3.5 ionising photons per hydrogen atom to reach a volume-averaged neutral fraction of  $x_{\text{HI}} = 0.01$ . The other colored curves represent three different ways of treating recombinations. The solid black curve uses the clumping factor of all IGM gas and computes the recombination term in Equation 10 as:  $C_b k_1 (10^4 \text{K}) x_{\text{HII}} \langle n_H \rangle$ , where  $x_{\text{HII}} = 1 - x_{\text{HI}}$  represents the ionised fraction. Down to  $z = 10$ , this crude treatment is already a significant improvement over the  $C = 0$  curve. Below  $z = 10$ , recombinations begin to win over the declining ionising emissivity, which is in turn driven by the assumed  $f_{\text{esc}}(z)$ ; the result is a multimodal reionisation history. The fact that the simulation does not, in fact, yield a multimodal reionisation history emphasizes the importance of allowing overdense regions to self-shield against the ion-



**Figure 6.** The dependence of the reionisation history on the definition of the clumping factor, as compared to our fiducial radiation hydrodynamic simulations. In photon-starved reionisation scenarios, a rapidly-growing clumping factor can lead to a double-reionisation history although this does not happen with our preferred definition of the clumping factor.

ising background. The dotted red and long-dashed magenta curves correspond to the definitions  $C_{\text{HII}}$  and  $C_{\text{HII},10^4\text{K}}$  from Section 4.1. In these models, the recombination term is  $Ck_1(10^4\text{K})x_{\text{HII}}^2(n_{\text{H}})$ . They are each a clear improvement over using  $C_b$ . Recall that the dotted red curve accounts for the order in which regions of different overdensities reionise while the long-dashed magenta curve additionally accounts for the impact of temperature fluctuations. Including this information accelerates reionisation by suppressing recombinations (Figure 3).

In detail, reionisation is more extended in the numerical model even when the clumping factor accounts for both the temperature and the ionisation fields (long-dashed magenta). We attribute this to the fact that the analytical model effectively assumes an infinite speed of light  $c$ . We have previously shown that artificially boosting  $c$  causes reionisation to complete earlier (Finlator et al. 2009b) although the effect was slightly weaker ( $\Delta z = 0.1$ ) than what is seen here ( $\Delta z \approx 0.9$ ). This delay is only a few times larger than the light travel time across the photon mean free path at this epoch ( $\sim 10$  Mpc; Furlanetto & Oh 2005), consistent with the possibility that light-travel time effects delay the completion of reionisation.

There are two additional effect that extend the final stages of reionisation in numerical models. The first results from the fact that the timescale for ionising an absorber varies inversely with its geometrical cross section. For example, an EUVB with amplitude  $\Gamma$  at the frequency corresponding to an absorption cross-section  $\sigma$  ionises a spherical absorber of uniform density  $n$  and radius  $r$  on a timescale  $\frac{2}{3}n\sigma r/\Gamma$  compared to  $1/\Gamma$  in the optically-thin case. Approx-

imating Lyman limit systems as uniform spheres in hydrostatic equilibrium at a temperature of  $10^4\text{K}$ , we find that systems with overdensities of 10 and 100 are ionised by a background  $\Gamma = 2 \times 10^{-13}\text{s}^{-1}$  on timescales of  $\approx 8$  and 24 Myr at  $z = 6$ , compared to  $\approx 0.2$  Myr in the optically thin case. This effect is omitted in analytical models but present in numerical simulations, although it probably delays the completion of reionisation by only  $\Delta z \approx 0.1$ .

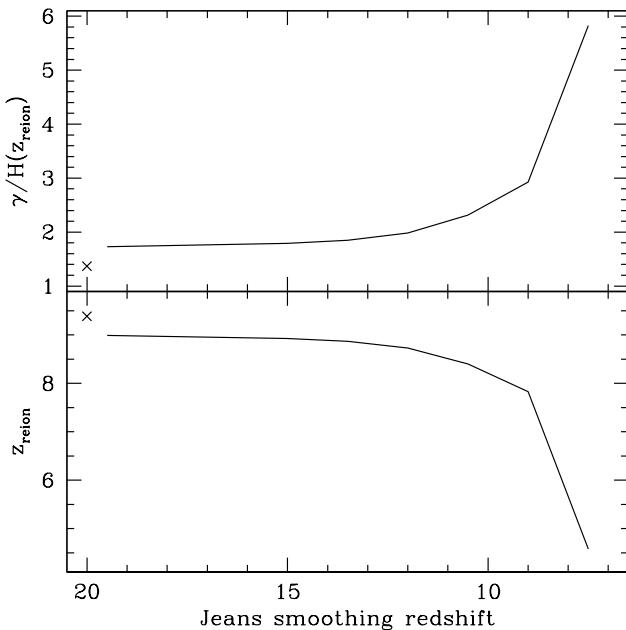
The second delay results from the fact that photons with high energy and small absorption cross section take longer to be absorbed than photons near the Lyman limit once the ionised volume fraction is large. This effect is missing from our simulations owing to our monochromatic radiation transport solver. However, it is not large: for  $z \geq 6$ , energies less than 54.4 eV, and neutral hydrogen fractions greater than 0.01, the delay is less than 20 Myr.

We may quantify the error in the reconstructed reionisation history using the optical depth to Thomson scattering  $\tau_{\text{es}}$ . We compute this for the simulations and the analytical models assuming that helium is singly-ionised with the same ionisation fraction as hydrogen down to  $z = 3$  (Ciardi et al. 2012; Friedrich et al. 2012), and doubly-ionised thereafter. In our simulation, the predicted value is 0.071. Ignoring recombinations yields  $\tau_{\text{es}} = 0.095$ , which confirms that galaxies could have dominated reionisation modulo uncertainties regarding the star formation efficiency of low-mass halos and the IGM recombination rate. Adopting realistic clumping factors measured directly from our simulation yields values for  $\tau_{\text{es}}$  that are between 0.062 (solid black) and 0.082 (long-dashed magenta). The discrepancy between the long-dashed magenta and dot-dashed green curves can be regarded as an estimate of the uncertainty in analytical reionisation calculations owing to complications such as light-travel time effects, shadowing, and source clustering; it is  $\approx 10\%$ . In short, analytical models of reionisation can reproduce the behavior of more realistic models reasonably well, with improvement possible if the adopted recombination rate accounts for the inhomogeneous temperature and ionisation fields. Our analytical modeling also confirms that collisional ionisations and the other complications that we have mentioned are subdominant.

We note that Gnedin (2000) previously performed an extensive comparison between his numerical radiation hydrodynamic simulations and analytical models, finding broad support for the basic assumptions underlying analytical models. Our results agree with his. The principal difference is that, in computing the volume-averaged recombination rate, their analysis did not distinguish between gas within galaxies and the IGM. This led to a much higher value for the clumping factor (100–200) than is appropriate for densities that are conventionally associated with the IGM (for example,  $\Delta < 1000$ ). Our simulations sidestep this complexity by using  $\Delta_{\text{lls}}$  as a physically-motivated model for the boundary between the IGM and the condensed gas whose absorptions are more appropriately attributed to  $f_{\text{esc}}$ .

## 6 APPLICATION II: THE PHOTON COST OF REIONISING FILAMENTS

While much has been learned about the photon cost of removing gas from minihalos (Haiman et al. 2001;



**Figure 7.** (bottom) The dependence of the overlap redshift on heating redshift: Delaying photoheating from  $z = 19.5$  to  $z = 7.5$  delays overlap by  $\Delta z = 3$ . (top) The reionisation photon budget per hydrogen atom as a function of heating redshift: Delaying photoheating can more than double the photon cost of reionisation. The crosses correspond to a homogeneous IGM ( $C = 1$ ).

Barkana & Loeb 2002; Shapiro et al. 2004; Iliev et al. 2005; Ciardi et al. 2006), the ability of moderately-overdense gas to delay reionisation has received less attention. This is potentially important for two reasons. First, analytical models may underestimate the reionisation photon budget by deriving the clumping factor from hydrodynamic simulations in which the IGM is effectively pre-heated. For example, it is possible to compute the clumping factor at high redshift by applying Equation 9 to the baryonic density PDF of Miralda-Escudé et al. (2000) (Wyithe et al. 2008; Pritchard et al. 2010). However, that PDF derives from simulations in which the diffuse IGM is already smoothed by an EUVB. In reality, there is a photon cost associated with smoothing the IGM because the initially cold IGM has a higher clumping factor (Figure 5). Ignoring this cost may cause models to underestimate the photon budget of reionisation. Alternatively, it is possible that the IGM is heated before reionisation is well under way. For example, an early X-ray background could suppress gas clumping prior to the onset of reionisation (Oh & Haiman 2003). In this case, filaments would absorb fewer ionising photons than in models (such as ours) in which the primary heating source is UV photons. In this section, we estimate the magnitude of this uncertainty by asking how many ionising photons are required to reionise filaments, or regions with overdensities of  $\Delta = 1\text{--}50$ .

Our approach involves decoupling reionisation from heating and determining how reionisation changes if the universe is heated at different redshifts. If the universe is heated at high redshift (such as  $z = 19.5$ ), then the photon cost of reionisation is smaller because the universe remains

smooth after being heated. By contrast, if the universe is heated at a lower redshift (such as  $z = 7.5$ ), then reionisation is more expensive because more of its gas has condensed into filaments by the time that reionisation is under way. By comparing these histories, we gain some insight (albeit not completely self-consistently) into the photon price that the universe pays for waiting to ionise its filaments; this is probably comparable to the actual photon cost of reionising moderately-overdense regions.

Pawlik et al. (2009) provide fitting functions for the baryonic clumping factor in simulations where an instantaneously imposed spatially-homogeneous EUVB is used to heat the IGM at a variety of redshifts. We have combined their fitting functions (choosing their  $C_{50}$ ) with the predicted ionising emissivity from our r6n256wWwRT16 simulation within our analytical model and evaluated how reionisation depends on the smoothing redshift. We show in the bottom panel of Figure 7 how the reionisation redshift (where  $x_{\text{HI}} \rightarrow 0.01$ ) varies. Delaying smoothing from  $z = 19.5$  to  $z = 7.5$  delays reionisation by  $\Delta z \approx 4.5$ . This relatively strong dependence owes to the extended reionisation history that is enforced by our choice of  $f_{\text{esc}}(z)$ . We have verified that extracting our emissivity history from simulations that assume a constant  $f_{\text{esc}} = 0.5$  yields a smaller delay of  $\Delta z = 1.1$  (not shown) because the growth rate of collapsed matter (such as stars) is much faster than the growth rate of moderately overdense structures (such as filaments).

The top panel maps the reionisation redshift into the number of ionising photons consumed per hydrogen atom. If the IGM is heated at  $z = 19.5$ , reionisation consumes 1.8 photons per hydrogen. Observations suggest that reionisation was well under way by  $z = 9\text{--}10$  (Pritchard et al. 2010; Kuhlen & Faucher-Giguère 2012). Adopting this as the photoheating redshift yields a total photon cost of 2.5–2.9 photons per hydrogen. Subtracting, we find that the universe pays a price of 0.7–1.1 photons per hydrogen for waiting to heat its filaments. Put differently, analytical models that assume that the universe is smoothed at a high redshift underestimate the total photon cost of reionisation by 0.7–1.1 photons per hydrogen atom.

These calculations may underestimate the total photon cost of reionisation because they do not account for gas that is condensed into minihalos, or halos with virial temperatures below  $10^4$  K ( $\approx 10^8 M_{\odot}$  for redshifts of 6–10). Prior to reionisation, much of the gas may have been condensed into halos with virial masses between  $10^6$  and  $10^8 M_{\odot}$ . The clumping factors of Pawlik et al. (2009) were derived from simulations that resolved halos at the upper end of this mass range with 100 particles, hence lower-mass systems are effectively missing. These halos can consume up to 5 photons per hydrogen atom (Shapiro et al. 2004), with the implication that the amount of gas clumping could have been much higher if the IGM temperature did not exceed 1000 K prior to reionisation.

On the other hand, an early X-ray background could deposit an early entropy floor due to the large mean free path of X-rays (Oh 2001; Venkatesan et al. 2001; Ricotti & Ostriker 2004; Madau et al. 2004). Such a background is widely predicted in many models of reionisation and could arise from X-ray binaries, inverse Compton scattering in supernova remnants, or early mini-quasars. This process significantly reduces gas clumping (Oh & Haiman

2003), and if it occurs well before the reionisation epoch then it could decrease the reionisation photon budget.

In short, while filaments may have absorbed  $\sim 1$  additional photon per hydrogen, the total number of photons required to complete reionisation remains uncertain owing to the unknown abundance of small absorbers. Future work incorporating X-ray heating and a wider dynamic range will be required in order to explore these processes.

## 7 DISCUSSION

Could galaxies have produced enough ionising photons to reionise the Universe? Could they have maintained the IGM's ionisation state at  $z = 6$ ? Two major uncertainties hamper efforts to connect observations of galaxies with measurements of the neutral hydrogen fraction in order to address these questions. The first is the fraction of ionising photons that escape galaxies' ISMs,  $f_{\text{esc}}$ , and the second is the recombination rate of the IGM, which is often parametrized by a clumping factor  $C$ . The observed population of galaxies can maintain an ionised IGM if the ratio  $C/f_{\text{esc}}$  lies within the range 1–10 (Finkelstein et al. 2010, Figure 15).

We previously showed that our simulations reproduce the observed UV luminosity function of galaxies and self-consistently complete reionisation by  $z = 6$  when we assume  $f_{\text{esc}} = 0.5$  (Finlator et al. 2011). While encouraging, those simulations suffered from two drawbacks. First, they overproduced the observed EUVB amplitude at  $z \leq 6$  and underproduced  $\tau_{\text{es}}$ . Second, the radiation transport solver's limited spatial resolution did not shield moderately-overdense regions from the EUVB, with the result that the predicted clumping factor of ionised gas was artificially high ( $\approx 12$  at  $z = 6$ ). Consequently, completing reionisation ( $x_{\text{HI,V}} < 0.01$ ) required up to 5 ionising photons per hydrogen atom.

In this work, we have improved on those calculations by adopting a subgrid self-shielding treatment and a redshift-dependent  $f_{\text{esc}}$ . Self-shielding limits the clumping factor and the reionisation photon budget by isolating overdense gas from the EUVB. Meanwhile, an evolving  $f_{\text{esc}}$  (1) reconciles the model with the observed  $\tau_{\text{es}}$  by increasing the electron fraction at early times; (2) heats the IGM at an earlier redshift, which suppresses gas clumping at later times; and (3) raises the Gunn-Peterson optical depth at  $z = 5$  into improved agreement with observations. Together, we find that these modifications suppress  $C$  to values in the range 2.7–3.3 at  $z = 6$  (the uncertainty comes from possibly overestimating the overdensity at which gas fully self-shields; see below). For our adopted  $f_{\text{esc}}(z)$ , this leads to  $C/f_{\text{esc}} \approx 25$ –30 at  $z = 6$ , slightly too high for currently-observed galaxy populations to maintain a completely ionised IGM. Our simulations overcome this barrier by resolving galaxies down to absolute magnitudes of  $M_{\text{UV}} = -15$ , which is 2–3 magnitudes fainter than current limits for observing galaxies directly but consistent with requirements inferred from other star formation tracers (Trenti et al. 2012).

Despite these improvements, there are a number of ways in which we can build on this work in the future. First, we use the global EUVB to compute  $\Delta_{\text{lls}}$ . Near sources, this underestimates the EUVB, leading to an overdensity threshold that is too low. In voids, it overestimates the EUVB, leading to too little self-shielding (see also Crociani et al. 2011).

It would be preferable to estimate the EUVB by averaging over regions that are a few times larger than the expected length scale of self-shielding systems in order to incorporate the EUVB's small-scale fluctuations into the ionisation field. This improvement would result in more-ionised overdensities and more-neutral voids (or in other words a more strongly inside-out reionisation topology), but the overall impact on  $C$  is difficult to predict.

Second, our results would benefit from an improved understanding of the IGM's temperature  $T$ . The IGM recombination rate is sensitive to  $T$  through the recombination coefficient  $k_2$  and through the minimum overdensity  $\Delta$  at which gas is neutral, which in turn is proportional to the self-shielding threshold  $\Delta_{\text{lls}}$ . For temperatures between  $10^{3.5}$ – $10^{5.5}$ ,  $k_2 \propto T^{-0.77}$ . Meanwhile,  $C(< \Delta_{\text{thr}}) \propto \Delta^{0.25}$  (Figure 5) and  $\Delta_{\text{lls}} \propto T^{0.3}$  (Equation 2), leading to a weak overall dependence of  $C(< \Delta_{\text{thr}}) \propto T^{0.075}$ . Combining these, we find that the overall IGM recombination rate varies as  $T^{-0.7}$ .

Uncertainty in  $T$  stems from the unknown latent heat of photoionisation and from resolution limitations. We have effectively assumed that reionisation heats gas to  $\approx 15,000$  K. While this assumption yields a mean IGM temperature that is consistent with observations (Section 2), the observations are themselves uncertain owing to the necessity to correct the measured IGM temperature for the influence of nearby quasars (Bolton et al. 2012). Assuming a (harder, softer) ionising background would lead to a (hotter, colder) post-reionisation IGM. Additionally, the limited spatial resolution of our radiation solver may cause our simulations to underestimate the post-reionisation gas temperature because ionisation fronts are artificially broadened, increasing the amount of time during which partially-ionised gas can cool through collisional excitation of neutral hydrogen (Miralda-Escudé & Rees 1994). The magnitude of this effect is likely small compared to the uncertain latent heat of photoionisation. Nevertheless, improved measurements of the IGM temperature at  $z > 5$  would constrain our model further.

Third, our simulations expose somewhat too much overdense gas to the EUVB, boosting the IGM recombination rate. For example, the threshold overdensity for self-shielding at  $10^4$  K and  $z = 6$  should be  $\Delta_{\text{lls}} = 25$  (Equation 2). If the neutral column density of self-shielded systems varies strongly with  $\Delta$ , then systems with overdensity more than a few times greater than  $\Delta_{\text{lls}}$  should be essentially neutral. For example, the high-resolution calculations of McQuinn et al. (2011) indicate that the neutral fraction should increase rapidly above  $\Delta \sim 100$  (their Figure 3). By contrast, the gas in our simulations remains ionised until  $\Delta \approx 200$  (Figure 5), which may be slightly too high. If so, then it owes largely to the fact that our simulations assume that all photons have an energy of 1.3 times the threshold energy for photoionisation  $h\nu_{\text{LL}}$  in order to photoheat the IGM. As we argued in the discussion of Figure 2, this assumption may artificially boost  $\Delta_{\text{lls}}$  by a factor of 1.7. Invoking once again the scaling  $C \propto \Delta_{\text{lls}}^{0.25}$  (Figure 5), we find that, if the amplitude of the ionising background were unchanged but the radiation field were dominated by photons with energy nearer to  $h\nu_{\text{LL}}$ , then the clumping factor of ionised gas would be lower by  $\approx 12\%$ . Combining this uncertainty with the range of values inferred from our fiducial and high-

resolution simulations (r6wWwRT16 and r6wWwRT32), we conclude that the clumping factor at  $z = 6$  lies between 2.7–3.3.

Finally, our limited cosmological volume forces us to adopt a rather dramatic evolution for  $f_{\text{esc}}(z)$ . We do this because our predicted  $\tau_{\text{es}}$  and  $\tau_{\alpha}$  are not converged with respect to simulation volume or mass resolution, as is well known for simulation volumes of this size (Gnedin & Fan 2006; Bolton & Becker 2009). This is not to say that the need for a decreasing  $f_{\text{esc}}(z)$  is purely a resolution limitation; empirical arguments strongly support such a trend and in fact our adopted relation lies within the observationally-inferred range (Kuhlen & Faucher-Giguère 2012). However, with a larger volume, structure formation would begin sooner, allowing us to adopt a weaker redshift-dependence for  $f_{\text{esc}}(z)$  without compromising the agreement between the observed and predicted  $\tau_{\text{es}}$  and  $\tau_{\alpha}$ . This would change the predicted IGM temperature-density relationship as well as the overall gas density distribution. Both of these changes could modify the predicted IGM recombination rate, although the sign of the effect is difficult to predict.

## 8 SUMMARY

We have used a suite of cosmological radiation hydrodynamic simulations to explore the impact of outflows and an EUVB on the recombination rate of the reionisation-epoch IGM. Feedback processes modify the IGM’s gas density, ionisation, and temperature fields. We may illustrate their impact on the gas density distribution by considering the clumping factor  $C_{100} \equiv \Delta^2 \mathcal{P}(\Delta)$  at redshift  $z = 6$  and overdensity  $\Delta = 100$ ; the overall clumping factor  $C_b$  (which neglects the ionisation and temperature fields) scales with this value.

- In the absence of feedback, baryons follow dark matter, and  $C_{100} = 10.7$  as previously found in N-body simulations (Iliev et al. 2007; Raičević & Theuns 2011).
- Galactic outflows return star-forming gas into the IGM, boosting  $C_{100}$ . The increase is less than 30% for any EUVB amplitude (including zero).
- An ionising background pressurizes the diffuse IGM, suppressing  $C_{100}$  to 2.0.

In reality, the IGM’s volume-averaged recombination rate depends on its density, ionisation, and temperature fields. Taking this information into account one piece at a time within the context of our fiducial simulation, we find (still at  $z = 6$ ):

- The clumping factor of all baryons below our adopted threshold density for star formation is  $C_b = 16.4$ .
- Averaging only over the ionised gas yields a much lower value of  $C_{\text{H II}} = 4.9$  because overdense gas is self-shielded.
- Accounting for the temperature field further suppresses the recombination rate to  $C_{\text{H II}, 10^4 \text{ K}} = 3.3$  because gas in the recent aftermath of reionisation is somewhat hotter than the canonical  $10^4$  K.
- The temperature-corrected clumping factor of ionised regions increases from 0.79 at  $z = 15$  to 3.3 at  $z = 6$ . Equation 8 reproduces this evolution reasonably well throughout  $z = 15 \rightarrow 5$ .

- Our fiducial simulation may overestimate the minimum density of neutral gas. Correcting for this may suppress  $C_{\text{H II}, 10^4 \text{ K}}$  by an additional factor of 12%. Hence our most realistic estimate of the ionisation- and temperature-corrected clumping factor at  $z = 6$  is that it lies within the range 2.7–3.3.

We have constructed an analytical reionisation model and tested how well different definitions of the clumping factor reproduce our numerical simulation’s reionisation history. We find that the clumping factor averaged over all IGM baryons  $C_b$  significantly overestimates the IGM recombination rate because it treats self-shielded gas as if it were optically thin. Accounting for the ionisation field through  $C_{\text{H II}}$  improves agreement with the numerical simulation, although the recombination rate is still overestimated because the post-reionisation IGM is in reality slightly hotter than the canonical  $10^4$  K. By contrast, using  $C_{\text{H II}, 10^4 \text{ K}}$  to account additionally for the IGM temperature field yields a monotonic reionisation history whose  $\tau_{\text{es}}$  overestimates the numerical result by only 10%. We speculate that this small remaining difference owes at least partly to light-travel time effects that are missing from the analytical calculation although other effects may also contribute.

We have used our analytical model to estimate the photon cost of reionising filaments, or regions with  $\Delta = 1\text{--}50$ . With our simulated emissivity history, if some process (other than reionisation) heats the IGM at  $z = 19.5$  then the photon cost of reionisation is  $\gamma/H = 1.8$  photons per hydrogen atom. Delaying the redshift at which the IGM is heated from  $z = 19.5$  to  $z = 9.0$  increases  $\gamma/H$  to 2.9 because the clumping factor increases. The difference between these numbers indicates that the Universe pays a price of  $\approx 1$  photon per hydrogen in order to reionise its filaments.

The ionisation state of dense gas has observational implications that merit further study. Our simulations show that, in the presence of self-shielding, outflows can boost the amount of gas at densities corresponding to the circumgalactic medium ( $10 < \Delta < 1000$ ). If a significant fraction of this gas remains neutral, then it could play two important roles in modulating the growth of the EUVB. Within halos, it could constitute a significant absorbing column, modifying the fraction of ionising photons that escape the host halo. This effect could mimic a redshift-dependent  $f_{\text{esc}}$ , which seems to be required by observations (Inoue et al. 2006; Finlator et al. 2011; Haardt & Madau 2012; Kuhlen & Faucher-Giguère 2012). Additionally, gas that travels past the virial radius while retaining a self-shielded component could boost the abundance of Lyman limit systems, as suggested by van de Voort et al. (2012).

A complementary motivation for improving our understanding of self-shielded gas is that it can be observed directly in the form of damped Lyman  $\alpha$  systems and low-ionisation metal absorbers. The IGM’s ionisation field is a major theoretical uncertainty hampering efforts to interpret observations of low-ionisation metal ions. For example, Oppenheimer et al. (2009) studied low-ionisation metal absorbers through the use of homogeneous ionising backgrounds. While this approach may be appropriate for ions such as SiIV and CIV, it is less appropriate for low-ionisation species such as OI and SiII because it artificially ionises,

heats, and smooths gas that should in reality be self-shielded. Our current simulations allow dense gas to self-shield, hence they will yield improved predictions for the abundance of low-ionisation metal absorbers. This will reduce the theoretical parameter space and increase the power of current observations.

## ACKNOWLEDGEMENTS

We thank the referee for a thoughtful and detailed report that improved the draft. We thank M. McQuinn and M. Prescott for helpful conversations. KF thanks D. Godinez and J. Sivapalan for exploring the ability of galactic outflows to self-shield within our simulations. Our simulations were run on the University of Arizona's Xeon cluster. Support for this work was provided by the National Science Foundation through grant AST-0907998, the NASA Astrophysics Theory Program through grant NNG06GH98G, and grant number HST-AR-10647 from the SPACE TELESCOPE SCIENCE INSTITUTE, which is operated by AURA, Inc. under NASA contract NAS5-26555. Support for this work, part of the Spitzer Space Telescope Theoretical Research Program, was also provided by NASA through a contract issued by the Jet Propulsion Laboratory, California Institute of Technology under a contract with NASA. KF gratefully acknowledges support from NASA through Hubble Fellowship grant HF-51254.01 awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS 5-26555. SPO acknowledges support from NSF grant AST-0908480.

## REFERENCES

Abel, T., & Haehnelt, M. G. 1999, *ApJL*, 520, L13

Barkana, R., & Loeb, A. 1999, *ApJ*, 523, 54

Barkana, R., & Loeb, A. 2002, *ApJ*, 578, 1

Barkana, R., & Loeb, A. 2004, *ApJ*, 609, 474

Bolton, J. S., & Haehnelt, M. G. 2007, *MNRAS*, 382, 325

Bolton, J. S., & Becker, G. D. 2009, *MNRAS*, 398, L26

Bolton, J. S., Becker, G. D., Raskutti, S., et al. 2012, *MNRAS*, 419, 2880

Bouwens, R. J., Illingworth, G. D., Oesch, P. A., et al. 2011, *ApJ*, 737, 90

Ciardi, B., Scannapieco, E., Stoehr, F., et al. 2006, *MNRAS*, 366, 689

Ciardi, B., Bolton, J. S., Maselli, A., & Graziani, L. 2012, *MNRAS*, 423, 558

Crociani, D., Mesinger, A., Moscardini, L., & Furlanetto, S. 2011, *MNRAS*, 411, 289

Davé, R., Finlator, K., & Oppenheimer, B. D. 2006, *MNRAS*, 370, 273 (DFO06)

Dunlop, J. S., McLure, R. J., Robertson, B. E., et al. 2012, *MNRAS*, 420, 901

Efstathiou, G. 1992, *MNRAS*, 256, 43P

Eisenstein, D. J., & Hu, W. 1999, *ApJ*, 511, 5

Faucher-Giguère, C.-A., Lidz, A., Hernquist, L., & Zaldarriaga, M. 2008, *ApJ*, 688, 85

Fernandez, E. R., & Shull, J. M. 2011, *ApJ*, 731, 20

Fan, X., Strauss, M. A., Becker, R. H., et al. 2006, *AJ*, 132, 117

Finkelstein, S. L., Papovich, C., Giavalisco, M., et al. 2010, *ApJ*, 719, 1250

Finlator, K., Özel, F., & Davé, R. 2009, *MNRAS*, 393, 1090

Finlator, K., Özel, F., Davé, R., & Oppenheimer, B. D. 2009, *MNRAS*, 400, 1049

Finlator, K., Davé, R., & Oumzel, F. 2011, *ApJ*, 743, 169

Friedrich, M. M., Mellem, G., Iliev, I. T., & Shapiro, P. R. 2012, *MNRAS*, 421, 2232

Furlanetto, S. R., Zaldarriaga, M., & Hernquist, L. 2004, *ApJ*, 613, 1

Furlanetto, S. R., & Oh, S. P. 2005, *MNRAS*, 363, 1031

Furlanetto, S. R., & Oh, S. P. 2009, *ApJ*, 701, 94

Gnedin, N. Y., & Ostriker, J. P. 1997, *ApJ*, 486, 581

Gnedin, N. Y. 2000, *ApJ*, 542, 535

Gnedin, N. Y., & Fan, X. 2006, *ApJ*, 648, 1

Gnedin, N. Y., Kravtsov, A. V., & Chen, H.-W. 2008, *ApJ*, 672, 765

González, V., Labbé, I., Bouwens, R. J., Illingworth, G., Franx, M., & Kriek, M. 2011, *ApJL*, 735, L34

Grazian, A., et al. 2011, *A&A*, 532, A33

Haardt, F. & Madau, P. 2001, in Clusters of Galaxies and the High Redshift Universe Observed in X-rays, 64 (HM01)

Haardt, F., & Madau, P. 2012, *ApJ*, 746, 125

Haiman, Z., Abel, T., & Madau, P. 2001, *ApJ*, 551, 599

Hopkins, P. F., Quataert, E., & Murray, N. 2012, *MNRAS*, 421, 3522

Hui, L., & Gnedin, N. Y. 1997, *MNRAS*, 292, 27

Iliev, I. T., Shapiro, P. R., & Raga, A. C. 2005, *MNRAS*, 361, 405

Iliev, I. T., Mellem, G., Shapiro, P. R., & Pen, U.-L. 2007, *MNRAS*, 376, 534

Inoue, A. K., Iwata, I., & Deharveng, J.-M. 2006, *MNRAS*, 371, L1

Katz, N., Weinberg, D. H., & Hernquist, L. 1996, *ApJS*, 105, 19

Kennicutt, R. C. 1998, *ApJ*, 498, 541

Kohler, K., Gnedin, N. Y., & Hamilton, A. J. S. 2007, *ApJ*, 657, 15

Komatsu, E., Smith, K. M., Dunkley, J., et al. 2011, *ApJS*, 192, 18

Kuhlen, M., & Faucher-Giguère, C.-A. 2012, *MNRAS*, 423, 862

Mac Low, M.-M., & Ferrara, A. 1999, *ApJ*, 513, 142

Madau, P., Haardt, F., & Rees, M. J. 1999, *ApJ*, 514, 648

Madau, P., Rees, M. J., Volonteri, M., Haardt, F., & Oh, S. P. 2004, *ApJ*, 604, 484

Martin, C. L. 2005, *ApJ*, 621, 227

McLure, R. J., Dunlop, J. S., de Ravel, L., et al. 2011, *MNRAS*, 418, 2074

McQuinn, M., Lidz, A., Zahn, O., Dutta, S., Hernquist, L., & Zaldarriaga, M. 2007, *MNRAS*, 377, 1043

McQuinn, M., Lidz, A., Zaldarriaga, M., et al. 2009, *ApJ*, 694, 842

McQuinn, M., Oh, S. P., & Faucher-Giguère, C.-A. 2011, *ApJ*, 743, 82

Miralda-Escudé, J., & Rees, M. J. 1994, *MNRAS*, 266, 343

Miralda-Escudé, J., Haehnelt, M., & Rees, M. J. 2000, *ApJ*, 530, 1

Murray, N., Quataert, E., & Thompson, T. A. 2005, *ApJ*,

618, 569

Muñoz, J. A., & Loeb, A. 2011, *ApJ*, 729, 99

Nestor, D. B., Shapley, A. E., Steidel, C. C., & Siana, B. 2011, *ApJ*, 736, 18

Oesch, P. A., Bouwens, R. J., Illingworth, G. D., et al. 2012, arXiv:1201.0755

Oh, S. P. 2001, *ApJ*, 553, 499

Oh, S. P., & Haiman, Z. 2003, *MNRAS*, 346, 456

Oppenheimer, B. D., & Davé, R. 2008, *MNRAS*, 387, 577

Oppenheimer, B. D., Davé, R., & Finlator, K. 2009, *MNRAS*, 396, 729

Pawlik, A. H., & Schaye, J. 2009, *MNRAS*, 396, L46

Pawlik, A. H., Schaye, J., & van Scherpenzeel, E. 2009, *MNRAS*, 394, 1812

Powell, L. C., Slyz, A., & Devriendt, J. 2011, *MNRAS*, 414, 3671

Pritchard, J. R., Loeb, A., & Wyithe, J. S. B. 2010, *MNRAS*, 408, 57

Raičević, M., & Theuns, T. 2011, *MNRAS*, 412, L16

Ricotti, M., & Ostriker, J. P. 2004, *MNRAS*, 352, 547

Schaerer, D. 2003, *A&A*, 397, 527

Schaye, J. 2001, *ApJ*, 559, 507

Schaye, J. 2004, *ApJ*, 609, 667

Schaye, J., Dalla Vecchia, C., Booth, C. M., et al. 2010, *MNRAS*, 402, 1536

Shapiro, P. R., Iliev, I. T., & Raga, A. C. 2004, *MNRAS*, 348, 753

Shapley, A. E., Steidel, C. C., Pettini, M., Adelberger, K. L., & Erb, D. K. 2006, *ApJ*, 651, 688

Shull, J. M., Harness, A., Trenti, M., & Smith, B. D. 2012, *ApJ*, 747, 100

Siana, B., Teplitz, H. I., Ferguson, H. C., et al. 2010, *ApJ*, 723, 241

Springel, V., & Hernquist, L. 2003, *MNRAS*, 339, 312

Springel, V. 2005, *MNRAS*, 364, 1105

Sutherland, R. S. & Dopita, M. A. 1993, *ApJS*, 88, 253

Theuns, T., Mo, H. J., & Schaye, J. 2001, *MNRAS*, 321, 450

Trac, H., & Cen, R. 2007, *ApJ*, 671, 1

Trac, H., & Gnedin, N. Y. 2009, arXiv:0906.4348

Trenti, M., Stiavelli, M., Bouwens, R. J., et al. 2010, *ApJL*, 714, L202

Trenti, M., Perna, R., Levesque, E. M., Shull, J. M., & Stocke, J. T. 2012, *ApJL*, 749, L38

van de Voort, F., Schaye, J., Altay, G., & Theuns, T. 2012, *MNRAS*, 421, 2809

Venkatesan, A., Giroux, M. L., & Shull, J. M. 2001, *ApJ*, 563, 1

Weiner, B. J., Coil, A. L., Prochaska, J. X., et al. 2009, *ApJ*, 692, 187

Wong, W. Y., Moss, A., & Scott, D. 2008, *MNRAS*, 386, 1023

Wyithe, J. S. B., Bolton, J. S., & Haehnelt, M. G. 2008, *MNRAS*, 383, 691

Table 1. Our suite of simulations. The fiducial simulation is indicated in bold.

name	$L^a$	winds	RT grid	$M_{h,\min}/M_\odot^b$
r6n256nWnRT	6	no	–	$1.4 \times 10^8$
<b>r6n256nWwRT16</b>	<b>6</b>	<b>no</b>	<b>16<sup>3</sup></b>	<b><math>1.4 \times 10^8</math></b>
r6n256wWnRT	6	yes	–	$1.4 \times 10^8$
r6n256wWwRT16	6	yes	16 <sup>3</sup>	$1.4 \times 10^8$
r6n256wWwRT32	6	yes	32 <sup>3</sup>	$1.4 \times 10^8$
r6n128wWwRT16	6	yes	16 <sup>3</sup>	$1.1 \times 10^9$
r3n128wWwRT8	3	yes	8 <sup>3</sup>	$1.4 \times 10^8$
r6n256nWHD01	6	no	–	$1.4 \times 10^8$

<sup>a</sup>in comoving  $h^{-1}$ Mpc<sup>b</sup>virial mass of a halo with 100 dark matter and SPH particles.

## APPENDIX A: APPROXIMATIONS TO THE EDDINGTON TENSOR

Our method of solving the radiative transport equation involves using a time-independent ray-tracing calculation to update the Eddington tensors in a way that tracks the evolving emissivity and opacity fields without becoming too computationally burdensome (Finlator et al. 2009a). The radiation transfer simulations in Table 1 use two approximations beyond those described in Finlator et al. (2009a).

First, we smooth each element of the Eddington tensor field with a 27-cell cubical tophat filter. This is necessary because the Eddington tensor can change rapidly at the positions of ionisation fronts, which can in turn cause numerical errors when the radiation field is updated. We introduced this approach in Finlator et al. (2009a) but have not previously used it within the context of cosmological simulations. We stress that it does not materially degrade our spatial resolution because the smoothing occurs over the same length scale as the finite-difference stencil that we use to discretize the moments of the radiation transport equation.

Second, we update each cell’s Eddington tensor less frequently as it becomes optically thin. Here, the assumption is that a region’s Eddington tensor evolves rapidly as an ionisation front sweeps over it and more slowly afterwards. We implement this idea as follows: The code only updates a cell’s Eddington tensor if the cell’s photon number density changes by more than a factor  $f_{\mathcal{J}}$ . Whereas previously we set  $f_{\mathcal{J}} = 0.05$ , we now allow  $f_{\mathcal{J}}$  to depend on the cell’s optical depth  $\tau = \chi \Delta x_{\text{RT}}$ , where  $\chi$  and  $\Delta x_{\text{RT}}$  are the local opacity and the cell size, respectively. In particular,  $f_{\mathcal{J}} = 0.05$  if  $\tau \geq 1$  and it grows linearly from 0.05 to 1 as  $\tau$  decreases from 1 to 0:

$$f_{\mathcal{J}} = \begin{cases} 0.05 & \tau \geq 1 \\ 1.0 - 0.95\tau & \tau < 1 \end{cases} \quad (\text{A1})$$

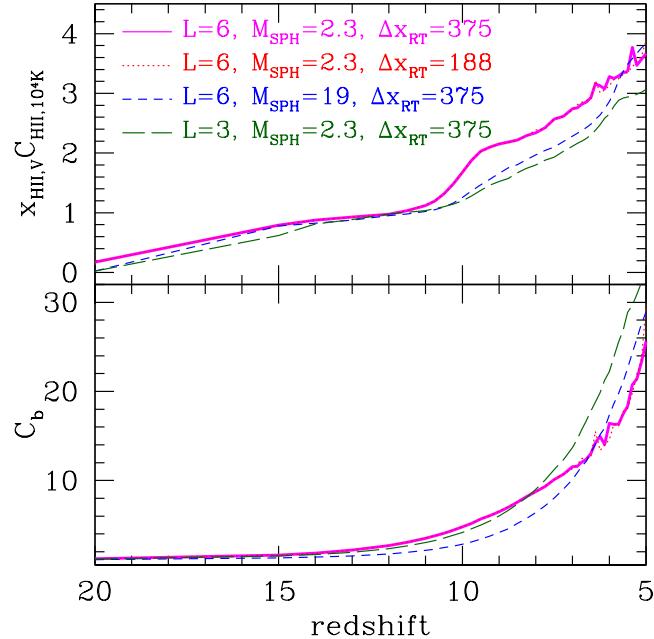
With this assumption, the Eddington tensor is updated frequently before the cell is reionised and less frequently once the radiation field has been established.

## APPENDIX B: RESOLUTION LIMITATIONS

In this section, we assess the sensitivity of our predicted clumping factors to three different flavors of resolution limitation: The simulation volume, baryonic mass resolution, and the spatial resolution of our radiation transport solver. We evaluate convergence in Figure B1 by comparing the baryonic and observational temperature-corrected clumping factors from different simulations with different numerical parameters but the same physical assumptions.

### B1 Simulation Volume

Our limited simulation volume could impact our results by modifying the relative contribution of voids and overdensities to the overall clumping factor (Bolton & Becker 2009; Raičević & Theuns 2011). By delaying the growth of structure (Barkana & Loeb 2004), small volumes also change the dependence of ionisation and temperature on overdensity at fixed redshift. Note that we have mitigated some of these problems by tuning  $f_{\text{esc}}(z)$  so that our simulations match a variety of observations (although not all observations; Section 2), with the result that the simulations are *representative* even if they are not *converged* (see also Gnedin & Fan 2006). Nonetheless, it is of interest to test our sensitivity to our limited volume. The solid magenta and long-dashed green curves in Figure B1 correspond to simulations that subtend different cosmological volumes at the same physical resolution, and their predicted temperature-corrected clumping factors of ionised gas (top panel) are reasonably close. In detail, the simulation with the larger volume (solid magenta) predicts a slightly higher recombination rate at all times, with the gap growing to 20–30% by  $z = 5$ . This difference owes entirely to the fact that, in a larger volume, reionisation begins sooner, hence at any given redshift it has proceeded farther from the voids into the overdense regions. This is especially true in our simulations because our self-shielding treatment effectively enforces an outside-in reionisation



**Figure B1.** The impact of various resolution limitations on the temperature-corrected clumping factor of ionised gas. The simulation volumes' side lengths  $L$  are in  $h^{-1}\text{Mpc}$ , the SPH particle masses are in  $10^5 M_\odot$ , and the radiation transport grid sizes are in comoving  $h^{-1}\text{kpc}$ . Resolution limitations are generally weak, indicating that our fiducial simulation volume (solid magenta) is reasonably converged numerically.

topology on small scales. In other words, for the same dependence of ionisation fraction and temperature on overdensity, the predicted recombination rates would be converged.

Turning to the bottom panel, we see that the small-volume simulation has a higher volume-averaged recombination rate  $C_b$ . This may also be attributed to the delayed progress of reionisation: At any given redshift, photoheating has smoothed the density fluctuations up to a lower density threshold in the small-volume simulation because the ionising background is weaker. Consequently, the clumping factor averaged over all IGM baryons (including the self-shielded ones) is higher.

This test does not account for fluctuations in the density field that are on larger scales than our largest simulation volume, which is comparable to the typical size scale of ionised regions when the Universe is 50% ionised (Furlanetto & Oh 2005). Similarly, it does not account for fluctuations in the EUVB on large ( $> 10h^{-1}\text{Mpc}$ ) scales, which can in turn source large-scale fluctuations in the ionisation field (Crociani et al. 2011). Future work at higher dynamic range will be required to assess these limitations.

## B2 Baryonic Mass Resolution

The mass resolution of our SPH discretization dictates the amount of dense structure that can form, which in turn impacts the recombination rate by regulating the progress of reionisation. The solid magenta and short-dashed blue curves in the top panel of Figure B1 illustrate these effects. The result of reducing the mass resolution at constant volume is qualitatively similar to reducing the simulation volume at constant resolution because both delay reionisation. In particular, the curve corresponding to the higher-resolution simulation (solid magenta) climbs systematically above its lower-resolution counterpart below  $z = 13$  because reionisation has proceeded into regions of higher overdensity at a given redshift, boosting the mean recombination rate. Following reionisation ( $z < 7$ ), the recombination rate in the lower-resolution simulation rejoins the prediction from the high-resolution simulation because the predicted gas density distributions and ionising backgrounds are similar. Turning to the clumping factor of all baryons  $C_b$  in the bottom panel, we see that the simulation with lower mass resolution has lower clumping prior to reionisation because the amount of matter that collapses into filaments and halos is lower at lower mass resolution. Following reionisation,  $C_b$  is higher at lower mass resolution because its weaker ionising background does not penetrate as far into overdense regions.

It is important to note that varying the baryonic mass resolution changes the predicted clumping even if the spatial resolution of the radiation solver is unchanged. This illustrates the benefit of solving the cooling and ionisation equations on the SPH particles rather than on the RT grid. On the other hand, the change is not large despite the order of magnitude increase in resolution, indicating that the requirements for numerically resolving the clumping factor are not strict and that our simulations are therefore reasonably converged.

### B3 Radiation Field Resolution

A central benefit of modeling a spatially-resolved EUVB is the ability to resolve self-shielding within overdense regions. Once reionisation is well under way, the boundaries of shelf-shielded regions dominate the IGM opacity and the clumping factor. Increasing the radiation field's resolution treats self-shielding more accurately, preventing overdense regions from becoming photoionised and smoothed. While our subgrid self-shielding prescription prevents the most overdense regions from being photoionised, which limits the clumping factor, our radiation solver's grid size is sufficiently large that it could still ionise regions near the self-shielding threshold too efficiently; the result would overestimate gas clumping. Additionally, modeling the radiation field with a low spatial resolution artificially broadens ionisation fronts, extending the time during which partially-ionised gas cools through collisional excitation of neutral hydrogen and underestimating the post-reionisation temperature. Our fiducial simulation's spatial resolution corresponds to an optical depth of 2 at the Lyman limit and the mean density, hence this could affect our results (Miralda-Escudé & Rees 1994) In order to explore these possibilities, we compare the solid magenta and dotted red curves in Figure B1, which show how the clumping factor evolves when the radiation field is discretized using our fiducial and twice the fiducial resolutions, respectively. The predictions are nearly indistinguishable, indicating that the clumping factors are insensitive to the resolution of our radiation transport solver. This resolution convergence is a clear demonstration of the power of our subgrid self-shielding prescription.

There are other numerical effects that we cannot consider directly using convergence tests. First, our use of a monochromatic radiation solver likely results in underestimating the IGM temperature, particularly in voids, because it does not capture spectral hardening (Abel & Haehnelt 1999) or heating owing to absorption of high-energy photons by helium atoms (Ciardi et al. 2012). This would slightly suppress the IGM recombination rate. On the other hand, the inability of our hydrodynamic solver to resolve minihalos means that our simulations may underestimate the IGM recombination rate (Haiman et al. 2001; Shapiro et al. 2004; Iliev et al. 2007; Ciardi et al. 2006). Resolving the question of whether numerical limitations cause us to over- or underestimate the clumping factor will require simulations with significantly higher dynamic range.